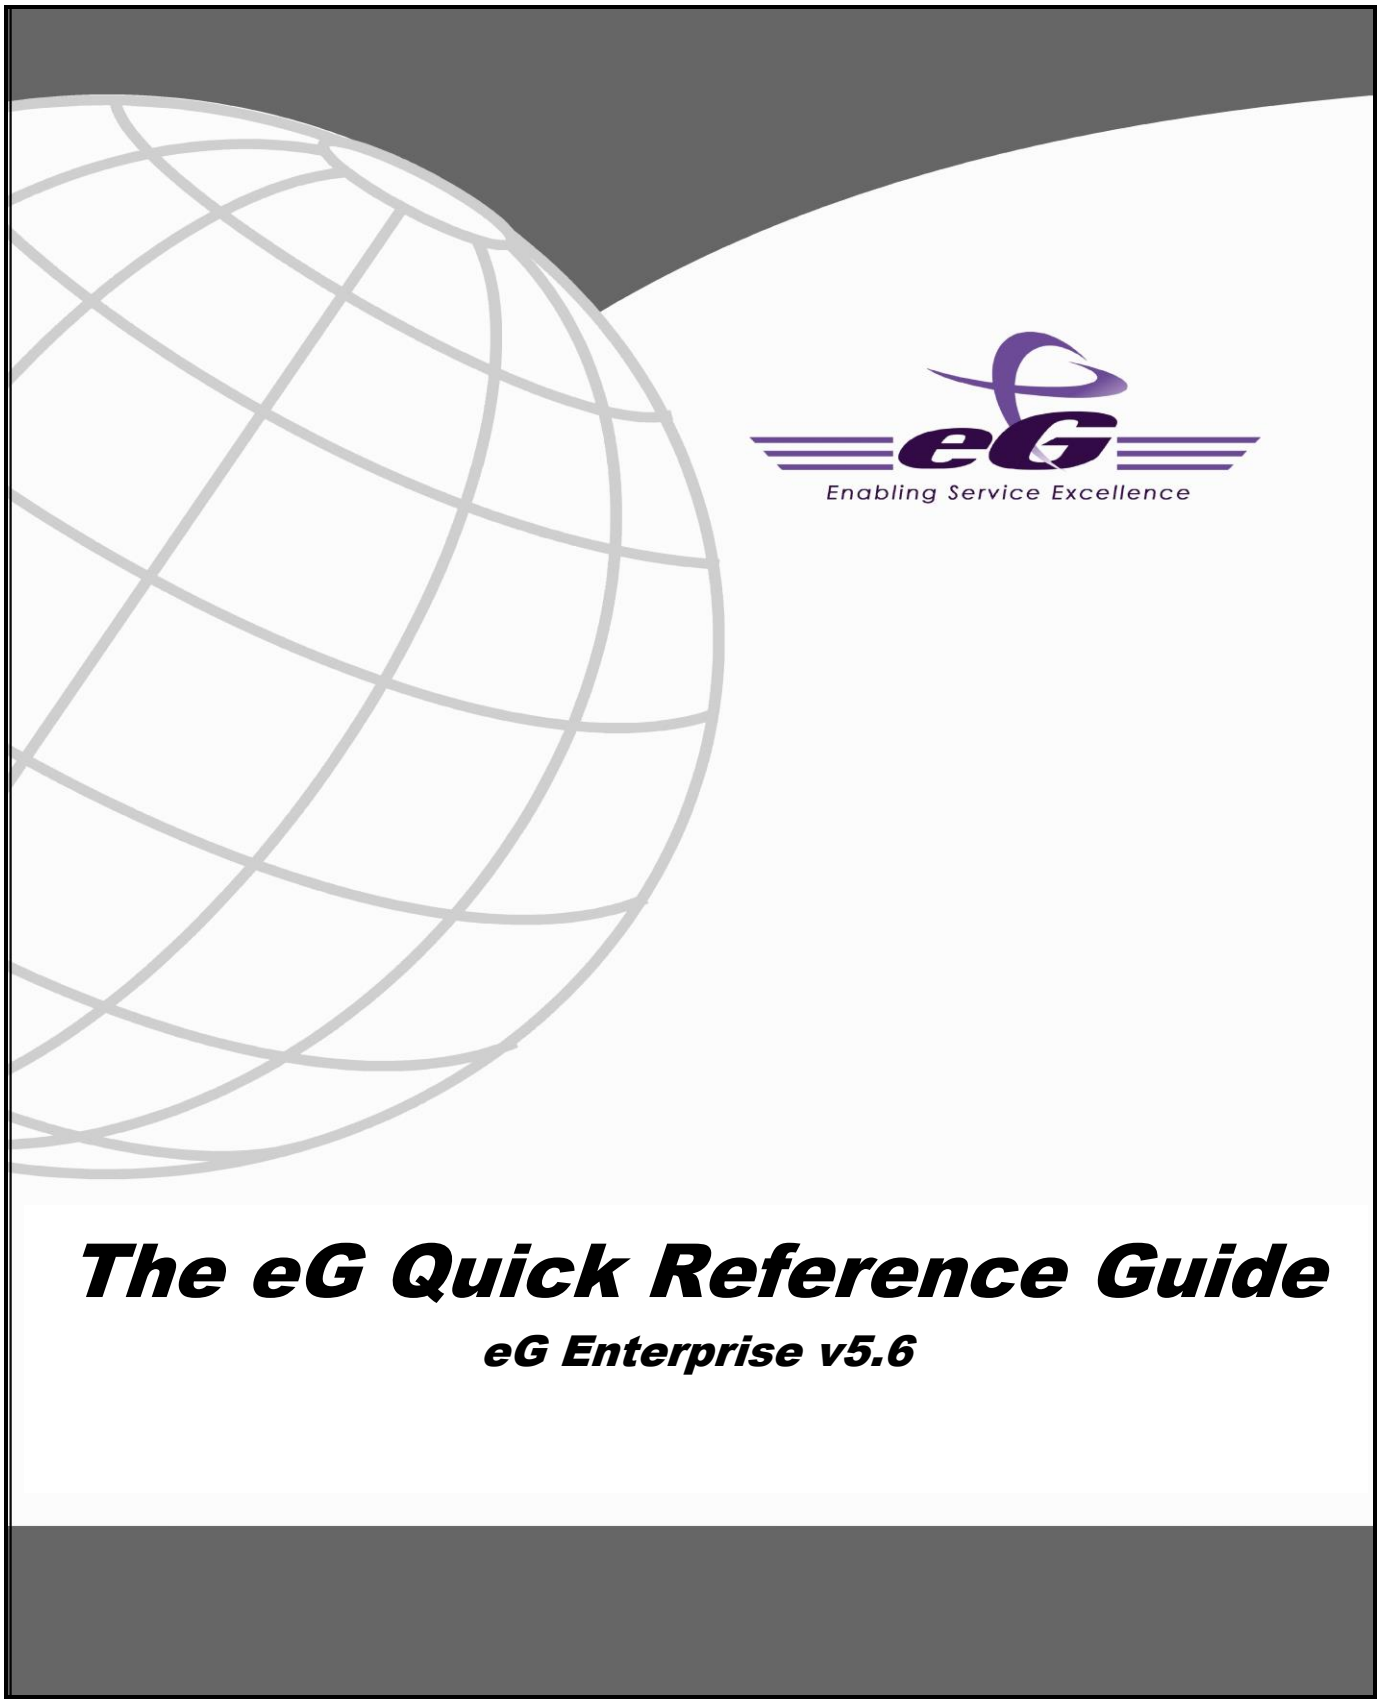

The eG Quick Reference Guide

eG Enterprise v5.6

Restricted Rights Legend

The information contained in this document is confidential and subject to change without notice. No part of this document may be reproduced or disclosed to others without the prior permission of eG Innovations Inc. eG Innovations, Inc. makes no warranty of any kind with regard to the software and documentation, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose.

Trademarks

Microsoft Windows, Windows NT, Windows 2000, and Windows 2003 are either registered trademarks or trademarks of Microsoft Corporation in United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Copyright

© 2013 eG Innovations, Inc. All rights reserved.

The copyright in this document belongs to eG Innovations, Inc. Complying with all applicable copyright laws is the responsibility of the user.

Purpose of the Manual

This manual provides answers to some of the most frequent queries that users have regarding the eG Enterprise Suite of products.

Table of Contents

INTRODUCTION TO MONITORING 1

ABOUT THE EG ENTERPRISE SUITE 3

SYSTEM ARCHITECTURE 8

3.1 Manager 8
3.2 Agents 8
3.3 Manager-Agent Communication 10
3.4 Database 11

LICENSING POLICY OF THE EG ENTERPRISE SUITE 12

INSTALLING AND CONFIGURING THE EG ENTERPRISE SUITE 16

5.1 Prerequisites for the Manager 16
5.2 Installing the Manager 17
5.3 Configuring the Manager 18
5.4 Starting the Manager 19
5.5 Prerequisites for the Agent 19
5.6 Installing the Agent 20
5.7 Starting the Agent 20

ADMINISTERING THE EG ENTERPRISE SUITE 22

6.1 Detailed Diagnosis Configuration 35
6.2 Supermanager Configuration 36
6.3 Configuring Reports 36
6.4 Maintenance Policies 39
6.5 Discovering Components 40
6.6 Managing/Unmanaging Components 43
6.7 Adding Components 44
6.8 Internal Agent Assignment 46
6.9 Configuring Zones 47
6.10 Configuring Groups 48
6.11 Configuring the Topology 49
6.12 Configuring Services 52
6.13 Configuring Transactions 54
6.14 Configuring Tests 54
6.15 Configuring Thresholds 55
6.16 Configuring Alarm Policies 58
6.17 Viewing the Agent Status 59
6.18 Manager Redundancy 60
6.19 Audit Logging 62

INFRASTRUCTURE MONITORING USING THE EG ENTERPRISE SUITE 63

CONFIGURATION MANAGEMENT 91

CONCLUSIONS 94

Chapter 1

Introduction to Monitoring

This chapter provides an overview of monitoring IT infrastructures

- **What is monitoring and how can it help me? Do I need a monitoring solution?**

  A monitoring solution is critical for on-going monitoring of an IT infrastructure. A few of the ways in which on-going monitoring can help include:

  - Early detection of problems
  - Knowledge of the normal usage of your IT infrastructure, so that anomalies are detected when they happen (e.g., a sudden surge in usage, a hack attack, etc.)
  - Analysis of the load-handling capacity of the different IT infrastructure components and to plan for capacity upgradation if required

- **What are the IT infrastructure components that will need to be monitored?**

  The performance of a service depends upon various components - the performance of the network, the server hardware, the server's operating system, the individual application platforms (e.g., middleware) in use, and the applications involved.

- **I have a network monitoring solution. Isn't this enough?**

  Network monitoring solutions are optimized to monitor and manage key network elements such as routers, hubs, and switches, to monitor network connectivity of the different servers, and to track usage of the different network segments. But an IT infrastructure is a combination of network, system and application elements. All of these layers are critical for efficient functioning of your infrastructure and a problem in one of these layers will have a cascading effect with a potential to bring your complete infrastructure down. So it is critical that you monitor all these key elements of your infrastructure. Since network-monitoring solutions do not have the capability to interface with and monitor systems and individual applications, end-to-end application monitoring requires specialized monitoring solutions. By having a network monitoring solution you can only look at the health of your backbone network, but the systems and applications on top of it could be failing without your knowledge.

- **Do I need multiple monitoring solutions?**

  eG products offer integrated network, system, and application monitoring. Hence, if required, eG products can be used as the only monitoring solution in an IT infrastructure.

- **What are the characteristics I should look for in a monitoring solution?**

  An ideal monitoring solution should be able to:

  Perform integrated monitoring (i.e., network, system, and application monitoring) of your IT infrastructure

  Provide actionable information, not raw data and thereby minimize the amount of human intervention necessary to operate it

  Automatically differentiate between the cause and effect of problems when they happen

  Proactively detect and alert operators of problems before these problems affect the user experience

- **Is correlation important for a monitoring solution?**

  System administrators are often confronted with a flood of alarms when a single problem occurs in their environment. Identifying the exact source of a problem requires a great deal of expertise and experience. Moreover, problem identification also takes time. Correlation capabilities included in a monitoring solution can dramatically reduce the time and effort necessary for problem diagnosis. To perform correlation effectively, the monitoring solution should be able to correlate between the network, system, and application layers.

- **We are currently using homegrown scripts for monitoring our infrastructure. Will this not suffice?**

  A whole lot of expertise is necessary in putting together a comprehensive end-to-end monitoring solution. The main drawbacks of homegrown scripts are:

  - Homegrown scripts are typically constructed in an ad-hoc manner, to address problems that have arisen in the past. Because they are put together in an ad-hoc manner, these scripts may not be complete, comprehensive, and efficient.
  - The cost of producing documentation, maintenance, and periodic upgradations of homegrown scripts is often very expensive.

Chapter 2

About the eG Enterprise Suite

This chapter gives a basic idea about the eG Enterprise Suite and its capabilities.

- **How will eG's products benefit me?**

  eG's products can enable customers to:

  - Proactively monitor and manage their infrastructure thereby offering the best performance for users
  - Instantly diagnose and fix problems, thereby minimizing down times and preserving their online revenues
  - Lower IT support costs by making monitoring simple and more effective
  - For service providers (hosting providers, Internet Data Centers, and Managed Service Providers), eG products can open up new revenue streams. The eG Enterprise Suite's virtual manager architecture allows service providers to offer revenue-producing monitoring and optimization services to their customers.

- **How does the eG Enterprise Suite differ from other monitoring solutions available in the market?**

  The primary differentiators of eG products relative to competing solutions are:

  Single click problem diagnosis capability based on a patent-pending correlation technology that correlates across network, system, and application layers to identify the root-cause of problems. Unlike many other solutions, this capability is available out-of-the-box and does not require elaborate set-up and customization

  - Real-time monitoring of web servers, web sites, and even individual web transactions, without using explicit, expensive logging of requests
  - Automatic thresholding of collected metrics so as to simplify configuration and set-up
  - Pay-per-use subscriber management model
  - Personalized, real-time views for individual users
  - Simple "One and Only" agent licensing that makes eG a cost-effective solution

➢ **What are the servers that the eG Enterprise Suite can manage?**

eG products are capable of monitoring various network, system, and application elements in an integrated manner. The table below summarizes the monitoring capabilities of the eG Enterprise Suite.

Component Type	Component Brand
Web servers	Apache, Microsoft IIS, IBM HTTP Server, Oracle HTTP Server, Sun Java Web Server
Web application servers	WebLogic, ColdFusion, Sun Java Application server, Microsoft transaction server, WebSphere, SilverStream, Jrun, Orion, Tomcat, Oracle 9i OC4J, Oracle Forms Servers, Borland Enterprise Servers (BES), JBoss, Domino application server, GlassFish Enterprise Server
Database servers	Oracle, Oracle RAC, Microsoft SQL server, DB2 UDB, DB2 DPF, Sybase, MySQL, SQL clusters, Backup SQL, Intersystems Cache, PostGre SQL, Oracle RAC, DB2 DPF
Network devices	Cisco routers, Cisco Catalyst switches, Baystack hub, Cisco VPN, Network nodes, Local director, Juniper SA and DX Device, 3COM CoreBuilder switch, Big-IP/F5 Load Balancer, Brocade SAN switches, Alcatel switches, Generic Fibre Channel switches, Cisco SAN switches, Cisco CSS, Cisco ASA, F5 Big-IP Local Traffic Manager (LTM), Coyote Point Equalizer
Microsoft Applications	Active Directory, BizTalk server, Windows Internet Name Service (WINS), Terminal server, DHCP server, MS Print server, MS Proxy Server, MS File server, ISA Proxy server, Microsoft Dynamics AX, Windows clusters, MS SharePoint, FAST Search for SharePoint 2010
Firewalls	Check Point Firewall-1, Cisco Pix, Netscreen Firewall, FortiGate Firewall
Terminal servers	Citrix MetaFrame 1.8, XP server, Citrix XenApp server, Windows Terminal server
Other Citrix Products	Citrix Secure Gateway, Citrix Secure Ticketing Authority, Citrix Web Interface (NFuse), Citrix Access Gateway, Citrix Netscaler, Netscaler ADC
Email servers	Microsoft Exchange, Instant Messenger on the Exchange 2000 server, Lotus Domino R5, Sun Java Messaging, Qmail server, AsyncOS Mail
Backup servers	Veritas Backup Exec server
Messaging servers	MSMQ, IBM MQ, FioranoMQ, Novell Groupwise, Tibco EMS
SAP	SAP R/3 server, SAP Internet Transaction server (ITS), SAP Web Application server, MaxDB
Others, LDAP, DNS	LDAP and SunONE LDAP server, DNS and Windows DNS server, FTP, MTS, Event Logs, Tuxedo domain servers, Printers, Windows Domain Controller, NetApp filers and NetCache, SiteMinder Policy server, Radius server, COM+ server, Tcp server, ASP .NET server, Network File System on Solaris server and client, Network File System on Linux server and client, MS RAS server, MS Radius server, eDirectory server, Sun Ray server, HP Blade server, XUps, Endeca Search, Bluecoat AV, 2X Client Gateway, 2x Publishing Agent, 2X Terminal server, SunONE Directory Server, Cisco UCS manager, Egenera PAN Manager, Teratext Arbortext, Teratext Content Server, DoubleTake Availability, Marathon everRun

Virtual Infrastructures	VMware® ESX Servers 3/3.5/ESXi, Solaris Containers, Microsoft Virtual Server, Solaris LDoms, Citrix XenServer, VMware vCenter, Microsoft Hyper-V, AIX LPARs on IBM pSeries servers, IBM HMC server, Citrix Provisioning Server, Oracle VirtualBox, RHEV Server, RHEV manager
Connection Brokers	Virtual Desktop Manager, Leostream connection broker, Xen Desktop Broker, VMware View, Oracle VDI Broker
SAN Storage Devices	Hitachi AMS, Hitachi USP, HP EVA StorageWorks Array, IBM DS RAID Storage, EMC CLARiiON, Dell EqualLogic, NetApp USD
Other Operating Systems	Generic SNMP, Generic Netware, AS400 server, OpenVMS server
Siebel Enterprise	Siebel Web Server, Siebel Application Server, Siebel Gateway
Cloud Providers	Amazon EC2, vCloud Director

- **What is a generic server? What measures corresponding to a generic server does the eG Enterprise Suite report?**

  Basic monitoring of a host at the network and operating system levels can be performed by managing the host as a "Generic server" in the eG architecture. For a generic server, an eG agent monitors the server's CPU, memory, disk, and network statistics. In addition, if desired, specific application processes executing on the server can be monitored.

- **What platforms do the eG products support?**

  The following tables show the platforms on which the eG agent can run on:

Platform	Version
Solaris	7, 8, 9, or 10
Red Hat Linux	6.0
Windows	2000, 2003, 2008, XP, Vista, 7
AIX	4.3.3, 5.x, 6.1, 7
HP-UX	10 and above
Free BSD	5.4
Tru64	5.1

- **How can the eG Enterprise Suite's correlation capabilities benefit me?**

  IT infrastructures typically comprise of applications and network devices operating in conjunction to offer a service. Often, a problem with one component can ripple and affect all the other components. Consequently, operators are confronted with a flood of alarms reporting a variety of different failures. Sifting through the alarms and determining where the root-cause of the problem may lie is a time-consuming, laborious process. Moreover, a significant amount of expertise is essential to effectively troubleshoot problems. Correlation technology embedded in monitoring solutions can significantly reduce the burden on operators, by pointing them directly to the root-cause of problems. eG products include specialized heuristics incorporated into the manager software to enable problem diagnosis at a single mouse click!

- **Can I monitor networks that are protected by firewalls using eG products?**

  Yes. The eG Enterprise Suite uses a truly web-based architecture - all communications between the agents and the managers happen through HTTP. The port used for all eG communications is configurable at the time of setting up the eG manager. Moreover, since the eG architecture uses an agent-driven model (with the agents polling the manager), the eG manager can be hosted outside a firewall, with the agents being used within the firewall domain.

- **Can I monitor IT infrastructures using eG products from my web browser?**

  Yes. The eG manager is a complete web application and hence, both administration and monitoring using the eG products are enabled through a web browser.

- **I am already using a network monitoring system. Can the eG Enterprise Suite integrate it into my environment?**

  Yes. The eG manager can be configured to send SNMP traps to any network monitoring system console.

- **Does the eG Enterprise Suite use SNMP for measurements?**

  Many of eG's pre-defined tests make measurements using SNMP. Monitoring of routers, load balancers, many application servers etc is done using SNMP.

- **What are the stages involved in deploying the eG Enterprise Suite?**

  Deploying the eG Enterprise Suite in a business environment is a very simple task. The following are the different steps involved in deploying the eG Enterprise Suite:

  1. **Installation** of the eG manager and the agents. This stage mainly involves deployment of the software on the appropriate servers, creating user accounts, and setting up the directory structures.
  2. **Configuration** of the eG manager and the agents. In this stage, the environment is set up for the proper operation of eG and the manager and agent processes are started.
  3. **Administration** of the eG. At this stage, the user interacts with the eG manager through the eG manager user interface to determine where agents must be deployed, what tests these agents must run, how often the tests should run, etc.
  4. **Monitoring** using the eG Enterprise system. At this stage, using the user interface, users can monitor various aspects of their IT infrastructure.

Figure 2.2 depicts the various stages involved in deploying eG Enterprise Suite in a target environment.

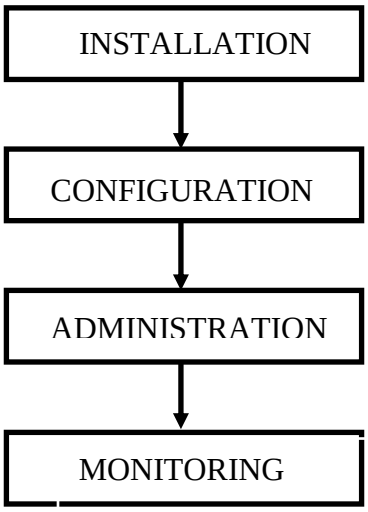

Figure 2.2: Stages involved in deploying eG Enterprise Suite

The following chapters will explain the details on the stages of deployment of eG Enterprise Suite and its use.

Chapter 3

System Architecture

A thorough understanding of the system architecture of the eG Enterprise Suite is essential for the user to deploy and use it effectively. This chapter elucidates the details of the eG Enterprise system architecture.

3.1 Manager

- **Does the eG database need to run on the same system as the manager?**

  It is not mandatory to host the eG database on the same system as the manager. In the event the eG database is hosted on a server other than the one hosting the eG manager, the firewall rules should also allow the manager-database communications.

- **Should the eG manager be deployed inside or outside a firewall?**

  The eG manager can reside either within or outside the firewall. While the Apache/Tomcat communications are internal to the eG manager, accesses from users and agents to the manager involve remote communication to and from the manager's web server port (7077 is the default port). If there are any firewalls used in the target environment, it is essential to ensure that the firewalls are configured to allow all communications to and from the web server component of the eG manager.

3.2 Agents

- **What is an agent? What are the types of agents that the eG Enterprise Suite supports?**

  Agents are software components deployed at various points in the IT infrastructure. Agents use different approaches for testing the target environment. The tests can be executed from locations external to the servers and network components that are responsible for the operation of the infrastructure. Agents that make such tests are called **external agents**. These agents take an external view of the target environment and indicate if the different services supported by the environment are functioning properly or not.

  Often external agents alone may not be sufficient to completely gauge the health of an IT infrastructure and to diagnose problems when they occur. For example, it may not be possible to measure the CPU utilization levels of a web server from an external location. To accommodate such situations, the eG Enterprise Suite uses **internal agents**. An internal agent runs on a server

that supports the IT infrastructure and monitors various aspects pertaining to the server (e.g., CPU, memory, and disk utilization, the processes executing on it, and the applications).

- **How many external agents does an eG manager support?**

  An eG manager has the capability to support multiple external agents.

- **What is a measurement?**

  The agents monitor the environment by running periodic **tests**. The outputs of the tests are called **measurements.** A measurement determines the state of a network / system / application / service element of the target environment. For example, a **Process test** reports the following measurements:

  - Number of processes of a specific type executing on a system.
  - The CPU utilization for these processes
  - The memory utilization for these processes

- **Does the eG Enterprise Suite provide separate agents for databases, web servers, application servers etc.?**

  Many network/system management platforms (e.g., BMC) provide separate agents for databases, web servers, application servers, etc. Having to install separate agents for the different servers poses a lot of overhead on the user. To simplify its installation and use, eG Enterprise Suite comprises of just a manager and an agent package. The agent package includes special-purpose tests for the different applications. Based on the discovery process, the eG manager determines which tests need to be deployed on each server.

- **How many agents can an eG manager support?**

  The number of agents that an eG manager can handle depends on various factors like:

  - The number and type of tests that each agent executes
  - The frequency with which each test runs
  - Whether the eG database is local or remote
  - The number of simultaneous monitor/admin users connected to the manager
  - How frequently the database is purged
  - The hardware on which the manager executes (the CPU and memory configuration of the system). No special hardware is required for the eG manager.

- **Does an agent listen on a port?**

  eG agents do not listen to any ports and therefore reduce security risk. Communications are always initiated from agents to manager. Statistical information/alerts are sent to the manager, which will respond with a timestamp of the last configuration change. New policies will be downloaded as required.

- **Do the agents run periodically?**

  An agent may comprise of multiple testers – one for each type of test that the eG Enterprise Suite supports. A tester executes tests one after the other. That is, if there are two Http tests that need to be run, the HttpTester will first run a test for a server A and after this, it will run another test for server B. A tester will attempt to run tests periodically. The tester will strive to keep the same average frequency as specified agents run pseudo-periodically – that is, they will attempt to run with the same average frequency, but the exact time at which a test is run will be varied. E.g., a test that has to be run once every 5 mins may be run after 4 mins once and after 6 mins the next time. By doing so, the eG agent attempts to catch performance phenomena that may happen with the same periodicity.

- **Can the eG Enterprise Suite use only external agents?**

  Yes. However, the diagnostic capability of the eG Enterprise Suite is critically dependent on internal agents. Without internal agents being deployed, an eG manager can detect infrastructure problems but may not be able to offer detailed diagnosis. An alternative to internal agents is the use of remote agents. A remote agent executes on a central server and collects OS/application metrics without needing an agent on each system.

3.3 Manager-Agent Communication

- **When there is network problem how does the eG agent and manager recover from it?**

  In the event of a network problem, the communication between the eG agent and the manager will be lost. As a result, the eG agent will be unable to report measurements to the manager. Under such circumstances, the eG agent stores the measurement information in local files. Once the network problem is fixed, the agent retrieves the information and passes it on to the manager. This ensures that measurements made by the agents are not lost even if there are intermittent manager / agent communication problems.

- **Is the eG agent intelligent enough to know that the manager is down? What happens if the agent is not able to talk to the manager?**

  Yes, the eG agent can identify that the manager is down. In this situation, as mentioned earlier, the agent retains all the measurement information in a local directory. Once the manager recovers, the agent retrieves all the information and sends it across to the manager.

- **What is the mechanism that is used for manager-agent communication?**

  The agent uses the push technology to put the information that it has collected into the manager. This way we are able to avoid opening of ports on the agent side, which many administrators and security experts are very concerned about.

- **What is the bandwidth required for agent-manager communication?**

  The factors controlling the required bandwidth are as follows

  (1) The number of agents reporting to the manager

(2) The frequency at which the tests are running

(3) The amount of information reported by the tests

In a typical environment, the manager/agent communication requires about 0.05 KB/s.

- **How does the eG agent communicate with the manager?**

  IT infrastructures typically include multiple demilitarized zones. From a security perspective, most IT infrastructure operators view SNMP and other proprietary protocols suspiciously, whereas they believe that HTTP/HPPTS traffic is not a serious security threat. Consequently, the eG Enterprise Suite uses HTTP/HTTPS for all communications between the manager and the agents.

- **Can the eG agent and manager be configured to use SSL (HTTPS)?**

  Yes. The eG agent and manager can be SSL-enabled, if so required. The advantage of enabling HTTPS support in the eG Enterprise Suite is that by default, the eG manager/agent communication is encrypted thereby preventing a third- party from snooping and decoding the data transmitted between the manager and agents. HTTPS support is particularly useful for remote monitoring across multiple locations, wherein the manager may be in a central location, and the agents at remote locations use the public Internet to communicate with the manager.

3.4 Database

- **What databases does the eG Enterprise Suite require for storing its data?**

  The eG Enterprise Suite requires Oracle database server (version 10G or higher) / Microsoft SQL Server (version 2005 / 2008 / 2012) for storing its data.

Chapter 4

Licensing Policy of the eG Enterprise Suite

This chapter clarifies the doubts on the licensing policy of the eG Enterprise Suite of monitoring products.

- **How is the licensing policy of eG Enterprise Suite different from that of the other monitoring tools?**

  Many monitoring systems have such complex licensing policies that just tracking the different licenses necessary to get a monitoring system operational itself is a burdensome task. Whereas, eG's licensing policy,

  - Simplifies the administration and use of eG products, by providing license management from a single, central location;
  - Requires minimal configuration and license management when adding new applications for monitoring;
  - Reduces the overheads involved in maintaining separate monitoring packages for individual applications.

- **What does eG's single agent licensing mean?**

  Many a time, system administrators and network operators have found that configuring and tracking the licenses required by a monitoring system itself can be even more time-consuming and laborious than managing their networks and servers directly! Many application monitoring tools enforce strict licensing on the software agents that are used to monitor servers and applications. Separate agent licenses need to be purchased based on various parameters like the deployment platform, the types and number of applications monitored, the number of CPUs, etc.

  Based on the market feedback, eG has introduced a powerful single agent licensing policy for its agents. As per this policy, a single eG agent can monitor all the applications executing on a server. Moreover, agent licenses are not tied to operating systems or node-locked, thereby allowing operators to pick and choose where they want to deploy the agents, and to even dynamically change the location of the agents. Its simple and cost-effective agent-licensing model makes eG an attractive solution for IT infrastructure monitoring.

- **How do I get a license for the eG products?**

  Once you get the demo software, you should send an email to license@eginnovations.com with the following information to obtain a valid license: IP address of the host on which you will run the manager, how many agents they want to use, and the address of their organization as well as a contact name.

- **How do I know if I have a valid eG license?**

  Issue the command "**/opt/egurkha/bin/viewCert /opt/egurkha/bin/license**" on Unix environments. On Windows environments, issue the command "**<target dir>\egurkha\bin\viewCert "<target dir>\egurkha\bin\license**". The output will provide information about the validity of the eG license. The same information is also available from the eG user interface.

- **Is it necessary to install a license on each server where an agent is to be installed?**

  No. The license is centrally managed by the eG manager. The manager is hooked to one IP address. Whereas, the agents have a restriction in the number and are not hooked to any IP addresses. The agents can be moved around. However, the number of running agents cannot exceed the number for which the license has been obtained.

- **Differentiate between Basic and Premium Monitors?**

  The eG license controls the total number of **Monitors** that can be used by an eG installation. The **Total Monitors** listing in the eG license indicates the total number of **basic** and **premium** monitors that the current installation of eG Enterprise is allowed to use, the total number of such monitors that are currently utilized, and the overall usage percentage.

  Basic Monitors are typically a combination of **basic agents** and **agentless basic components**. These concepts have been demystified below:

  - **Basic Agent** - A **basic** agent can be used to monitor only the operating system of a host and the processes running on it. To use a basic agent, the user must manage the host as any of the following: Generic server, Eventlog server, Windows Generic server, Linux, Solaris, AIX, HPUX, Windows server, MS File server, or MS Print server.
  - **Agentless Basic Components** - eG Enterprise supports both agent-based and agentless monitoring of .target components. For implementing *agentless monitoring*, the solution uses **Remote agents**. A remote agent implements agentless monitoring for one or more of the target servers/applications - i.e., without requiring an internal agent for the target server/application, the remote agent can collect critical statistics about the target. For more details regarding eG's *Agentless Monitoring* capability, refer to the *eG User Manual*.

  In eG parlance, hosts/applications that are monitored in an agentless manner are called **Agentless Components**. Typically, a Generic server, Windows Generic server, Linux, Solaris, AIX, HPUX, Windows server, MS File server, MS Print server, AS400 server, Netware server, Snmp Generic server, when managed by remote agents, will be counted as **Agentless Basic Components**.

  The eG license does not restrict the number of **Basic agents** or **Agentless Basic** components that can be configured in an eG installation - instead, it only restricts the total number of **Basic Monitors**. This imparts to customers the flexibility to decide how to use the basic monitors - i.e., they can decide how many of the **Basic Monitors** should be reserved for **Basic Agents** and how many for **Agentless Basic Components**. This way, administrators can move the licenses between agent-based and agentless components, without actually modifying the license.

  Premium Monitors are typically a combination of **premium agents, agentless premium components**, and **external agents**. These concepts have been demystified below:

  - **Premium Agents** - If any applications on the host (e.g., Web, email, DNS, etc.) have to be monitored, the internal/remote agent used for this purpose is counted as a **premium agent**.

- **Agentless Premium Components** - In eG parlance, hosts/applications that are monitored in an agentless manner are called **Agentless Components**. All applications that are monitored using a remote agent are called **Agentless Premium Components**.
- **External Agents** - The eG license restricts the number of external agents that can be configured in the target environment. Since each eG installation requires at least one external agent, the customer's license must allow for at least one.

The total number of **Premium Monitors** that an eG installation is allowed to use is automatically computed as the difference between the **Total Monitors** and **Basic Monitors**. For instance, if the eG license allows 20 monitors totally and 7 basic monitors, then 13 will be automatically set as the maximum number of premium monitors that can be configured in the environment. If required, the customer can even have 19 basic monitors, reserving 1 premium monitor for the 1 external agent that is a must for every eG installation.

Moreover, the eG license does not explicitly restrict the number of **Premium Agents** or **Agentless Premium** components that can be configured in a target environment - instead, it restricts the total number of **Premium Monitors**. This imparts to customers the flexibility to decide how to use these **Premium Monitors** - i.e., they can decide how many of the **Premium Monitors** should be reserved for **Premium Agents** and how many for **Agentless Premium** components. However, since the **Premium Monitors** restriction includes the ceiling on external agents as well, customers cannot exhaust their **Premium Monitors** license on premium agents and agentless premium components alone - in other words, if the eG license allows a total of 10 **Premium Monitors** and 3 **External Agents**, it implies that only 7 **Premium Monitors** can be consumed by premium agents and/or agentless premium components - 3 **Premium Monitors** will be reserved for external agents. This way, administrators can move the licenses between agent-based and agentless components, without actually modifying the license.

- **I need to monitor 5 VMware ESX servers in my environment. Each ESX server hosts 4 VMs. If I only want to monitor the host operating systems, and the status and resource usage of each of the guests, what would be the license requirement for both agent-based and agentless approaches?**

1. **Agent-based approach**

   1 Premium Monitor license for every ESX server to be monitored = 5 Premium Monitor licenses

   1 Premium Monitor license for the default external agent installed on the eG manager

   Total = 6 Premium Monitor licenses

2. **Agentless approach**

   No license is required for the remote agent that monitors all 5 ESX servers

   1 Premium Monitor license for every ESX server to be monitored by the remote agent = 5 Premium Monitor licenses

   1 Premium Monitor license for the default external agent installed on the eG manager

   Total = 6 Premium Monitor licenses

- **Does the eG license control the 'Agentless Monitoring' capability?**

  Yes. Agentless Monitoring is indeed a license controlled feature of the eG Enterprise Suite.

Chapter 5

Installing and Configuring the eG Enterprise Suite

This chapter delves into the details of installing and configuring the eG Enterprise Suite.

5.1 Prerequisites for the Manager

- **What are the hardware prerequisites for the eG manager?**

  a. A minimum of 2 GB RAM would be required for installing the eG manager on a 32-bit host; for a 64-bit host on the other hand, a minimum of 4 GB RAM would be required
  b. A minimum of 1 GB of disk space free

- **What are the software prerequisites for the eG manager?**

  1. **On Unix, the requirements are:**

     i. Solaris 10 (or higher) or Red Hat Enterprise Linux 5 (or higher), CentOS 5.2 (or higher)
     ii. JDK 1.6.0_10 (or above)
     iii. Oracle database server (version 10G or higher) / Microsoft SQL Server (version 2008 / 2012) for the eG database. The database can be installed on the same system as the eG manager, or it can be installed on a separate system. For implementations with 100 monitors or more, the database should ideally be hosted on a separate system. Both the eG manager and the eG database can be hosted on virtual machines or physical machines.
     iv. A valid eG license

  2. **On Windows, the requirements are:**

     i. JDK 1.6.0_10 and above
     ii. Windows 2003 server (OR) Windows 2008 server (OR) Windows XP workstation (OR) Windows 7 (OR) Windows 8 (OR) Windows 2012
     iii. Only systems with a static IP address (i.e. no DHCP address) should be used for installing the eG manager
     iv. Internet Explorer 9 or higher or Mozilla Firefox Version 16 or higher or Chrome.
     v. Oracle database server (version 10G or higher) / Microsoft SQL Server (version 2008/2012) for the eG database. The database can be installed on the same system as the eG manager, or it can be installed on a separate system. For

implementations with 100 monitors or more, the database should ideally be hosted on a separate system. Both the eG manager and the eG database can be hosted on virtual machines or physical machines.

v. A valid eG license

- **What are the database requirements for the eG manager?**

The database requirements for the eG manager are:

a. Must be able to provide the database administrator name and password to create a new user account for the eG user
b. Must be able to provide the default and temporary tablespace in which the eG user account is to be created
c. At least 100 MB of space available for the eG user. The exact size requirement depends on the target environment.

5.2 Installing the Manager

- **Do I need a dedicated server for the eG manager and agents?**

We recommend a dedicated server for the eG manager. The external agent executes on the same system as the manager. The system on which the eG manager and external agent executes is the **eG server**.

- **Can I have multiple eG managers?**

For scalability, an IT infrastructure can choose to deploy multiple eG managers. Alternatively, to ensure high availability, a single primary manager and multiple secondary managers can also be configured in the environment. However, an agent must work under the control of only one manager.

- **Can an existing user account be used to install the eG manager and agents?**

Yes. The installation process checks that the user account specified exists. If the user account exists, the installation uses the existing user account itself.

- **Can I use the same server for both the eG manager and the database?**

Yes, provided both the database server and the eG manager reside on the same host.

- **Can the manager and agent be installed using different user accounts?**

On different hosts, the manager and agent can be installed in different user accounts. However, on the same host, the manager and agents must be installed using the same user account. The installation directory should also be the same.

- **I am installing the eG manager. The manager install prompts for a database username/password. Do I need to create this user account in the database?**

  No. The install process will create the user account if it does not exist.

- **The manager install process prompts for data and temporary tablespaces when using an Oracle database as the backend. Do I need to manually create these tablespaces?**

  No. You can use existing data and temporary tablespaces. To find tablespaces, go to the sql prompt and type the command "**`select * from v$tablespace`**". This will report the existing tablespaces to you. Choose the data and temporary tablespaces from this.

- **I am installing an eG manager on a Windows box. My environment does not comprise of an Oracle or an MS SQL server. Is there an alternative database that I can use?**

  If you choose not to use an Oracle or an MS SQL server backend for a Windows 2003/Windows 2008/ Windows 7 / Windows 8 / Windows 2012 eG manager, then you can install and use the Microsoft SQL Server 2005/2008/2012 Express. The eG manager setup itself takes care of installing MS SQL 2005/2008/2012 Express. However, the respective executables need to be downloaded from the web and made available on the eG manager host. The download URLs and the steps for installation are detailed in the *eG Installation Guide*.

  However, note that, MS SQL Server 2005/2008/2012 Express can only serve as a temporary substitute for the MS SQL server, as it provides only limited storage and scalability capabilities. Owing to such constraints, it is strongly recommended that you restrict the usage of MS SQL Server 2005 Express to short-term monitoring of a relatively small number of components - to be precise, a maximum of 25 components. We also recommend that you acquire a licensed version of the MS SQL server as soon as possible, install the server, and migrate the eG database to it.

5.3 Configuring the Manager

- **What is the significance of the host name or the IP address that is being specified while configuring the eG manager?**

  While configuring the eG manager, if the user provides an IP address, the same should be used for connecting to the manager. If he/she specifies the host name, then only this host name should be used for connecting to the manager. Improper usage may result in login errors even when a valid user tries to login.

- **Can the eG manager be double-byte enabled? If so, when?**

  The eG Enterprise system provides users with the option to view and key in data in a language of their choice. Different users connecting to the same manager can view data in different languages. However, some languages such as Chinese, Japanese, and Korean, support a double-byte character set. To view data in the eG user interface in Chinese, Korean, or Japanese, the eG manager should be explicitly configured to display and process double-byte characters. In such a case, enable double-byte support for the eG manager by specifying **y**. On the other hand, for handling the character sets of other languages (example: French, German, Spanish, Portugese, etc.), the eG manager need not be double-byte enabled. At such times, enter **n** to disable double-byte support for the eG manager.

To double-byte enable an eG manager installation, simply type **y** (in case of a Unix manager), or click the **Yes** button (in case of a Windows manager), when manager setup prompts you for your confirmation to enable the usage of a double-byte character set. For the other steps to be followed to ensure complete double-byte support, refer to Chapter 4 of the *eG Installation Guide*.

5.4 Starting the Manager

- **When I start the eG manager, what processes can I expect to see?**

On Unix environments, look for the following processes:

a. A number of web server processes with a path of "`/opt/egurkha/manager/apache/bin/httpd`"
b. One tomcat process (search for a pattern "java –Xbootclasspath" in the ps output)
c. The eG manager's recovery process (eGmon)

On Windows environments, look for the following processes:

a. eGmon (manager recovery process)
b. eGurkhaTomcat (core manager process)

5.5 Prerequisites for the Agent

- **What are the hardware pre-requisites for the eG agents?**

a. At least 256 MB RAM
b. At least 1 GB of disk space free

- **What are the software pre-requisites for the eG agents?**

1. **On Solaris, the requirements are:**
   i. Solaris 7 or higher (use "uname –a" to check the OS version)
2. **On Linux, the requirements are:**
   i. Red Hat Linux 3 or higher (use "uname –a" to check the OS version) or openSUSE v11 (or above)
3. **On CentOS, the requirements are:**
   i. CentOS v5.2 (or above)
4. **On AIX, the requirements are:**
   i. AIX 4.3.3 (or higher)
5. **On HP-UX, the requirements are:**

i. HP-UX 10 or higher

6. **On Free BSD, the requirements are:**

   i. Free BSD 5.4

7. **On Tru64, the requirements are:**

   i. Tru64 5.1

8. **On Windows, the requirements are:**

   i. Windows 2000 server with Service Pack 4 (OR) Windows 2003 server (OR) Windows 2008 server (OR) Windows 2000 Professional workstation (OR) Windows XP workstation (OR) Windows Vista (OR) Windows 7 (OR) Windows 8 (OR) Windows 2012

5.6 Installing the Agent

- **Can multiple agents execute on the same host?**

  No. Only a single eG agent should execute on a host.

- **Does the agent need to be installed in a system to discover the components in it?**

  No. An agent is not required to be installed on a system in order to discover the components (applications) within it as the eG Enterprise suite follows selective port discovery.

5.7 Starting the Agent

- **When I start the eG agent, what processes should I see?**

  On Unix environments, look for the following processes:

  - EgMainAgent – the core agent process
  - A script named eGAgentMon that periodically monitors the agent and restarts it if the agent ever fails

  On Windows environments, look for the following processes:

  - eGurkhaAgent (core agent process)
  - eGAgentMon (agent recovery process)

- **I start an agent without starting a manager? Will this work?**

  No. The agent waits for a maximum of one day for the manager to come up.

Chapter 6

Administering the eG Enterprise Suite

This chapter mainly focuses on many of the common questions that an administrator of the eG Enterprise system may have. Please refer to the *User Manuals* for detailed information.

- **How do I connect to the eG manager?**

  Users can connect to the eG manager using the URL **http://<eGManagerIP>:<eGManagerPort>** (use **https://<eGManagerIP>:<eGManagerPort>**, if SSL-enabled).

- **How do I login to the Enterprise system?**

  The user will get to see the login window as soon as a connection is established with the eG manager from where he/she can access the eG Enterprise system. The eG Enterprise system is predefined with a default administrator account with a login "admin" and password "admin". Specify the same in the space provided and click the **Authenticate** button.

- **What if I forget my password while logging in?**

  Simply click on the **Forgot Password** button in the login page. In the page that appears next, provide the **Username** for which you have forgotten the password and click the **Get Password** button. The password will then be emailed to the mail ID that corresponds to the specified **Username**.

- **What is the first page that appears in the eG user interface, when you login as "admin"?**

  The eG Admin Home page appears.

- **What do the various sections of the Admin Home page represent?**

  This page enables the administrator to understand, at a glance, the status of the eG monitoring system. The page reveals the following information:

  - The first section of the page provides the **AGENT SUMMARY**. This section displays the total number of agents that have been configured for the environment. The number of agents of each type (Basic and Premium) that have been configured, and the number of agents that are currently running/not running are also indicated. Clicking on a bar corresponding to Premium or Basic agents will lead you to **STATUS** page that lists all the agents of the corresponding type, and their current status.
  - Below the **AGENT SUMMARY** is the **USER SUMMARY**. This section indicates the number of **Expired** user accounts, user accounts which are **Nearing Expiry** (i.e. the accounts that will expire within 7 days), and **Others**. The **Others** bar includes the number of users who will never expire (i.e. users for whom the **No Expiry** option has been set), and the number of users

who are not expected to expire within the next 7 days. Note that if any of the above-mentioned counters contain the value 0, then the corresponding bar will not be displayed. For instance, if there are no expired user accounts, then the **Expired** bar will not appear. Clicking on any of the bars in this section takes you to that will reveal the users who belong to the category 'clicked on', the zones, services, segments and components associated with each of the users, and the date of subscription expiry.

- For each component (i.e. network device or application) being monitored, eG includes a specialized model that dictates what tests must be run by the eG agent to monitor the component. Many of eG's tests are pre-configured i.e., do not require manual configuration. A few tests require explicit configuration. The third section in the **Admin Home** page provides the **UNCONFIGURED COMPONENT SUMMARY**. Besides providing the total number of unconfigured tests, this section graphically depicts the number and type of components for which tests remain to be configured. For more information on the unconfigured tests, click on the bar that corresponds to a component type. This will take you to the **LIST OF UNCONFIGURED TESTS** that provides the complete list of tests requiring manual configuration for the chosen component type.
- Adjacent to the **AGENT SUMMARY** section is the **LICENSE USAGE SUMMARY**. The eG license governs a wide variety of factors such as the number and type of agents that the installation can support, the number of applications that can be monitored, etc. The information provided by the **LICENSE USAGE SUMMARY** helps the administrator assess whether the eG licenses are being effectively utilized. Clicking on the **< <Click here to get the license details> >** link will take you to the **LICENSE INFORMATION** page that provides the license details as well as its usage details.
- Finally, a **COMPONENTS AT-A-GLANCE** section is available, which indicates the component-types that are currently monitored by eG Enterprise, and the number of managed components per type. Clicking on the bar corresponding to a component-type leads you to the **ADD/MODIFY** page where more components of that type can be added for monitoring.

- **Who are the default users to the eG Enterprise system?**

admin and *supermonitor*

- **What is a user role?**

In large enterprises, the IT staff have clearly demarcated roles and responsibilities. The help desk staff are responsible for handling user complaints and their main concern when a user calls about a problem is to determine whether the user call pertains to a problem that the other operations staff are already working on. The domain experts and service managers are responsible for the early detection, diagnosis and fixing of problems with the networks, servers, applications, and services they control. While the domain experts are interested in the detailed performance metrics relating to the IT infrastructure, the executive managers are interested in high-level service level reports that detail if the IT infrastructure is meeting the service expectation of their users. To support these varying requirements of the IT operations staff, the eG Enterprise Suite supports different user roles. The user roles define the rights and responsibilities that any user of the eG Enterprise system has. Each user in the eG Enterprise system is assigned to a user role.

- **Explain the default user roles of the eG Enterprise system.**

  - **Admin:** Users who are assigned administrative rights become the super-users of the system. Such users can choose what hardware and application servers are to be monitored by the

system, where the agents should be executed to monitor the hosted environment, what tests these agents should run, how often these tests should be executed, and can view the status of the entire monitored infrastructure. The administrative user also has the rights to add, delete, and modify user roles and individual user profiles. The default *admin* user is assigned the **Admin** role only.

- **Monitor:** Monitor users have restricted access to the eG Enterprise Suite. Each monitor user is associated with an email address to which alarms will be forwarded. The user's profile also includes information regarding the customer's alarm preferences – whether alarms have to be forwarded in text or HTML mode, whether a complete list of alarms has to be generated each time a new alarm is added, or whether the new alarm alone should be sent via email, etc. Each monitor user is associated with a subscription period. The eG Enterprise Suite allows the monitor users to access the system until this period only.
- **A Supermonitor** user has an unrestricted view of the monitored infrastructure. He/she can receive alarms pertaining to the whole infrastructure that has been configured by the administrative user. The default *supermonitor* user is assigned the **Supermonitor** role only.
- **AlarmViewer**: This role is ideal for help desk personnel. The users vested with **AlarmViewer** permissions can login to the monitor interface, and perform the following functions:
    - View the details of alarms associated with the specific components and services assigned to them
    - Provide feedback on fixes for the alarms
    - View feedback history
    - Change their profile

  Like **Monitor** users, users with this role can only monitor the components assigned to them.
- **SuperAlarmViewer**: Users with the **SuperAlarmViewer** role have all the privileges of the **AlarmViewer** role. In addition, users with the **SuperAlarmViewer** role have access to all the components being monitored.
- **ServerAdmin**: The users who have been assigned the **ServerAdmin** role have all the administrative rights of an **Admin** user, except **the right to user management**. Similarly, like a **Supermonitor** user, a **ServerAdmin** user can monitor the complete environment, and even change his/her profile.
- **MonitorNoConfig**: The users who have been assigned the **MonitorNoConfig** role will have access to the eG monitoring and reporting interfaces only, and not the eG **Configuration Management** interface.
- **SupermonitorNoConfig**: Users with **SupermonitorNoConfig** privileges will have unrestricted access to the monitoring and reporting consoles only - such users will not be able to access the configuration management console, even if the eG license enables this capability.

➢ **I created a user role and provided 'Limited' COMPONENT ACCESS to that role. Once this was done, I noticed that the 'Admin' section of the page grayed out all options within, except 3 options. Why did this happen and what does it mean?**

In any monitored environment typically, administrators alone have the right to make configuration changes using the eG administrative interface. Monitor users on the other hand have no access to the administration console. In large enterprises, multiple distinct administration teams may use the same eG Enterprise manager for their monitoring. These teams would require

the ability to configure the monitoring for the servers they operate. To address such environments, eG Enterprise included the capability to configure users with limited administration rights. For instance, a separate role can be created to allow monitor users with just the permissions to configure tests that should be executed on their servers, to change the thresholds that can be applied for monitoring their servers, and/or to apply maintenance policies on the servers assigned to them. This is why, as soon as the **Limited** option is chosen, all the check boxes except the **Agent Test Config**, **Agent Threshold Config**, and **Maintenance Policy Config** check boxes are grayed out in the **Admin** section. This implies that user roles with **Limited** component access can only perform one/more of the following administrative functions:

- Configuring tests pertaining to the components assigned to them
- Configuring the thresholds related to the components under their monitoring purview
- Suppressing the alerts related to the components assigned to them by configuring maintenance policies

On the other hand, if the **Complete** option is chosen, it implies that the user role has access to all the monitored elements in the infrastructure, and can be granted any administrative/monitoring privilege as the administrator deems fit.

- **Can I add a new user to the environment?**

Yes. As soon as an administrator successfully logs in, the first page to appear will be the Admin Home page. From the **Users** menu, select the **Add** option, which will enable the administrator to add a new admin or monitor user.

- **Can I change passwords for the default users?**

Yes. The default users of the eG Enterprise Suite are Admin and Supermonitor. The Admin user can change his password using the **Change Password** option present in the **Users** menu. An administrator can change the password of a Supermonitor using the **Modify User** option present in the same menu. .

- **What is alarm escalation? Can I escalate alarms?**

The eG Enterprise system can be configured to automatically escalate persistent problems to the next managerial level, so that the higher authority can initiate immediate action. When a user who receives an email/SMS alert of an issue is unable to resolve the reported issue for a pre-configured period of time, then the eG Enterprise system immediately forwards the email/SMS alert to a comma-separated list of mail IDs/mobile numbers specified against the **Level 1** field in the **Escalation Mail ID** section. You can, if you so desire, define additional support levels by clicking on the '+' sign that appears at the end of the **Level 1** text box. This way, issues that remain unresolved even at **Level 1** will be escalated to **Level 2** and so on. You can create upto a maximum of 5 escalation levels. To delete a newly added level, click on the '-' sign at the end of the corresponding **Level** text box.

Note:

Alarm escalation will work only if you configure the following:

- The duration beyond which the eG Enterprise system needs to escalate a problem to the next level
- The alarm priorities to be escalated

Both these parameters can be configured using the **ALARM ESCALATION** section in the **MAIL ALERT CONFIGURATION** page that appears when the Alerts -> Mail Settings -> Alerts menu sequence is followed.

- **Can I delete users from the eG Enterprise system?**

  Yes. The **Delete Users** option present in the **Users** menu will enable the administrator to remove an existing user from the eG Enterprise system.

- **Can I enable the remote control capability for specific users?**

  Yes. Using the **ADD USER** screen that appears upon clicking the **Add User** option in the **Users** menu, you can enable this capability for a specific user. This remote control capability allows monitor/supermonitor users to remotely manage and control servers from a web browser itself. From the browser, a monitor/supermonitor can execute commands on a monitored component, run some diagnosis, and initiate corrective actions to fix any problems.

- **When and why do I need to specify a Mail ID/Mobile No. while creating a new user?**

  The eG manager is capable of alerting users as and when problems occur. The alarms are classified into critical, major, and minor. When a mail id/mobile no. is associated with a user, then the eG Enterprise system will forward the alarms to the specified mail id(s)/mobile no(s). When multiple mail IDs are specified, an administrator can specify which mail address(es) need to be in the To: field of the mail alarm and which ones should be in the Cc: and Bcc: fields. If a mobile number(s) is specified, then a compact alarm report, ideal for a mobile phone console is generated. However, note that eG alarms will be forwarded to a mobile phone only if an **eG SMS Manager** has been installed in the network, and the eG manager has been configured to work with the SMS manager.

- **I want the email alert to a specific user to include the detailed diagnostics as well. Is this possible?**

  Yes. You can configure a user profile to additionally receive the detailed diagnosis (if any) associated with a problem measure in the email alerts. For this, just set the **Include detailed diagnosis in mail alerts** flag in the **ADD USER** page to **Yes**.

- **Besides problem information and detailed diagnosis, what more can an email alert contain?**

  Specific user profiles can be configured to receive additional measure details in email alerts. Such additions can be any of the following: a graph of the problem measure plotted for the last 1 hour (OR) the data plotted in a 1-hour measure graph.

- **Can I specify the alarm priorities that a user can view in the eG monitor console's Alarm window?**

  Yes. By selecting specific alarm priorities from the **ALARM DISPLAY** field in the **ADD USER** page, an administrator can configure the alarm priorities that a user can view while monitoring the infrastructure.

- **What is the significance of the 'Time zone' setting in the ADD USER page?**

  By default, all alerts generated by the eG Enterprise system are based on the eG manager's time settings. However, in an infrastructure that spans multiple geographies, users who are responsible for the proper functioning of servers in a particular geography (or time zone) may want to receive email alerts pertaining to those servers in their local time zone. To ensure this, you now have the option of configuring a time zone for a new user. The **Time Zone** list in displays a wide variety of time zones to choose from. By default, the manager's time zone is displayed here, indicating that, by default, email alerts are generated based on the manager's time settings. However, in case of target environments that are spread across multiple time zones, you may want to associate a different **Time Zone** with every user, so that email alerts sent to a user report problems based on that user's time settings, and not the manager's. In such a case, select the required **Time Zone** from the list. Once such a user is created, all subsequent email alerts that the user receives will report problems based on that user's **Time Zone** settings only.

- **Can I start and stop the manager via the user interface?**

  No. The user cannot start or stop the manager via the eG user interface.

- **Can I change profile for default users?**

  Yes. The profile of an existing eG user can be modified using the **Modify** option in the **Users** menu.

- **Can I configure a user profile with the right to delete alarms?**

  Yes. To achieve this, just set the **Allow Delete** flag in the **ADD USER** page to **Yes.**

- **After enabling the 'Remote control' capability for a user, I noticed that a 'Remote command execution' filed appears in the ADD USER page. What is this all about?**

  Using the **Remote command execution** filed you can indicate whether the new user is authorized to execute any command remotely, or is only allowed to choose from a pre-configured list of commands. For more details, refer to the *eG User Manual*.

- **What is alarm acknowledgement? How do I enable this capability for a user?**

  Optionally, specific users can be configured to **acknowledge** an alarm displayed in the eG monitor interface. By acknowledging an alarm, a user can indicate to other users that the issue raised by

an alarm is being attended to. In fact, if need be, the user can even propose a course of action using this interface. In such a case, a user with Admin or Supermonitor privileges (roles) can edit the acknowledgement by providing their own comments/suggestions on the proposed action. The acknowledgement thus works in three ways:

- Ensures that multiple members of the administrative staff do not unnecessarily invest their time and effort in resolving a single issue;
- Serves as a healthy forum for discussing and identifying permanent cures for persistent performance ills;
- Indicates to other users the status of an alarm

To enable the alarm acknowledgement capability for the new user, select the **Yes** option from the **Allow alarm acknowledgement** section in the **ADD USER** page.

- **The ADD USER page now allows me the flexibility to create a Local user or a Domain user. What is this all about?**

The eG administrative interface provides administrators with a wide variety of options to manage user information. Be it user creation, modification, deletion, or simply viewing user information, any type of user-related activity can be performed quickly and easily using the eG administrative console. Some target environments however, use the Active Directory server as the central repository for user information, and for authenticating domain logins. In such situations, if both the eG manager and the Active Directory consoles were to be used together for managing and validating the same set of users, it is bound to result in chaos! Therefore, to provide administrators with a single, central interface for efficiently managing users, eG Enterprise integrates seamlessly with the Active Directory server. The first step towards implementing this integration is the creation of a domain. Use the Users -> Configure Domains menu sequence to configure the domain, which is managed by the Active Directory (AD) server. Subsequent to domain creation, if you attempt to **Create New** user using this page, you will be prompted to indicate the **User authentication** mode that applies to the new user. If you are creating a domain user, whose requests are to be authenticated by the Active Directory, then select the **Domain** option. If you are creating a user who is local to the eG Enterprise system, and whose login requests are to be authenticated by the eG database, select the **Local** option. Upon choosing the **Domain** option, you will be prompted to perform the following tasks:

- Select the **Domain** to which the new user belongs. The domain that you created using the Users -> Configure Domains menu sequence will be listed in the **Domain** list.
- Then, set the **Operation** flag to **User**, specify the ID of the new user in the **User ID** text box, and click the **Validate** button. When this is done, the eG manager immediately connects to the Active Directory server and verifies whether the user is a valid domain user or not. If the user is not a valid user, then an error message to that effect appears. On the other hand, if the user is indeed a valid domain user, then the eG manager allows you to proceed with the user creation. However, you cannot provide a password for the domain user. This is because, the credentials of the domain user are configured in and maintained by the Active Directory server; eG Enterprise therefore, will neither reveal nor allow you to modify the password of the domain user, thus ensuring data integrity. Moreover, subsequently, when you log into the eG management console as a domain user, you will have to make sure that you prefix the user name with the domain name in the format: <<*domainName*>>/<<*Username*>> (or <<*domainName*>>\<<*Username*>>). Every time a domain user logs into the eG Enterprise system, the login will be authenticated by the Active Directory server that manages the users in that domain.

- **What is the purpose of the 'Operation' flag in the ADD USER page?**

  Using the **Operation** flag, you can indicate whether the profile being created is for a domain user or a domain group.

- **How do I create a domain group in the eG Enterprise system?**

  To create a domain user group, then, do the following:

  - Set the **Operation** flag to **Group**.
  - Next, select the **Domain** to which the group belongs. The domains that you created using the Users -> Configure Domains menu sequence will be listed in the **Domain** list.
  - Soon after a **Domain** is chosen, the **Group Name** list will be automatically populated with all the user groups that are pre-configured in the selected **Domain**. Pick the **Group Name** that is to be registered with the eG Enterprise system. All users who are part of this AD group will now be allowed access to the eG Enterprise system. The rights and privileges (eg., role, expiry date, email/SMS alert settings, alarm acknowledgement/deletion rights, etc.) defined for the chosen group will govern all users who belong to that group. This saves administrators the trouble of defining separate profiles for each domain user in a group.
  - **Note that the group is not associated with any 'password'**. This implies that while a *group* itself cannot login to the eG management console, a user who belongs to the group can login using the credentials defined for him/her in the AD server. At the time of login, the group user should provide his/her name in the format: <*DomainName*>\<*UserName*>. Everytime a group user logs into the eG management console, the solution automatically connects to the AD server to validate the login.

    Note:

    - If a domain user group is registered with the eG Enterprise system, and a profile is later created in eG for a particular domain user in that group, then, when that user logs into the eG management console, the user-level settings will override the group-level settings.
    - If a domain user belongs to more than one AD group that is created in the eG Enterprise system, then, when that user logs in, the solution provides him/her with a list of domain groups to choose from. Selecting a group from the list enables the user to automatically inherit the access rights and monitoring scope defined for that group.

- **I work only during the nights, and therefore, prefer to receive email alerts only during that time. Can eG Enterprise be configured to facilitate this?**

  Yes. By configuring shift periods for email/SMS/escalation alerts, you can ensure that these alerts are received only during specified time periods.

- **How to configure Email alerts to be sent out during shifts?**

  To configure shift periods for email alerts, do the following:

  - First enable the **Email alerts only during shift periods** flag, by setting it to **Yes** in the **ADD USER** page.

- Upon setting the flag to **Yes**, you will be required to specify the **Days** on which the user should receive email alerts; also, in the **Shifts** field alongside, you need to mention the specific time periods on the chosen **Days** at which the user should receive email alerts

- First, click on the **Calendar** control ( ) next to the **Days** field.
- From the **DAYS** list that pops out, which lists the days of the week, select the days on which email alerts need to be sent to the user.
- To choose more than one day from the list, select a day by clicking on the left mouse button, and then, with the **Ctrl** button on your keyboard pressed, click on another day to select it. Similarly, multiple days can be selected. To add your selection to the **Days** field, click the **Add** button in the **DAYS** list. You will thus return to the **ADD USER** page where the selected days will be listed against the **Days** field.
- Next, using the **Shifts** field, provide the specific time periods at which email alerts should be sent out to the user on the chosen days. For that, do the following:
- First, click on the **Calendar** control () next to the **Shifts** field. Doing so invokes the **SHIFTS** window, wherein you can specify a **From** time and **To** time for your shift. **Ensure that the shift timings correspond to the Time zone chosen for the user**.
- To provide an additional time slot, click on the circled '+' button at the end of the first row. Another row then comes up wherein you can provide one more time period. In this way, you can associate a maximum of 5 shift periods with the chosen **Days**.
- To remove a shift period from the **SHIFTS** window, simply click on the circled '-' button against the corresponding specification. Finally, to add these time periods to the **Shifts** field, click on the **Add** button in the **SHIFTS** window. You will thus return to the **ADD USER** page, where you can find the time period(s) that you specified appear against the **Shifts** field.
- To add another **Day-Shift** specification, just click on the circled '+' button at the end of the first row. Another row will then appear, where you can specify a few more **Days** and **Shifts**. This way, a number of **Day-Shift** specifications can be associated with a user.
- To delete a particular **Day-Shift** specification from, simply click on the circled '-' button.

➢ **Is there a limit to the number of Day-Shift specifications that can be associated with a user.**

Yes, but this number is configurable. It can be any number between or equal to 1 and 10. To configure this number, go to the **MAIL ALERT CONFIGURATION** page that appears when you follow the *Alerts -> Mail -> Alerts* menu sequence. In the **SHIFT PERIOD CONFIGURATION** section of this page, select the **Maximum number of Day-Shift combinations**.

➢ **How to configure shift periods for SMS alerts?**

Similar to email alerts, SMS alerts can also be configured to be sent out only during specified time periods on specific days of the week. The first step towards this is to enable the **SMS alerts only during shift periods** flag by selecting the **Yes** option in the add user page. Using the **Days** and **Shifts** fields that appear subsequently, you can configure one/more Day-Shift combinations in the same manner as discussed for email alerts**.**

- **I do not want the option to enable email/SMS/escalation alerts for shift periods to appear in the ADD USER page. What do I do to ensure this?**

  To ensure that the shift-related flags do not appear in the **ADD USER** page, perform the following:

  - Use the menu sequence *Alerts -> Mail Settings -> Alerts* to open the **MAIL ALERT PREFERENCES** page.
  - In the **Shift Period Configuration** section of the page, set the **Allow shift period configuration** flag to **No**. By default, this flag is set to **Yes**.
  - Finally, register the changes by clicking the **Update** button in that page.

- **I want a user to receive email alerts of only those issues that pertain to the Network layer. Can email alerts be so filtered?**

  By default, a user receives email/SMS alerts for all issues pertaining to all components assigned to him/her. In some circumstances, the user may not want to receive all of these alarms. For instance, in a large, multi-tier infrastructure, a user may be monitoring all the applications and network devices involved in supporting a business service. However, the user may have primary responsibility only for some of the components supporting the business service (e.g., a network administrator's primary responsibility is to monitor the network devices). In such cases, while the user may want to view the status of all the components of the business service, he/she may want to receive email or SMS alerts pertaining to specific components of the infrastructure alone (e.g., network devices).

  To enable such selective alerting, eG Enterprise provides administrators with the option to configure the eG manager to **not send out email/SMS alerts related to specific layers/components/component-types/tests for specific users**.

  By default, the ability to **filter mail/SMS alerts** is disabled. To enable it, do the following:

  - Use the menu sequence *Alerts -> Mail Settings -> Alerts* to open the **MAIL ALERT PREFERENCES** page.
  - In the **FILTER MAIL ALERTS** section of the page, set the **Allow mail/SMS filter configuration** flag to **Yes**.

  Once this is done, a **Configure mail/SMS filters** button will additionally appear in the page using which components/segments/services/zones/service groups are assigned to a user.

- **How can I modify the eG database settings?**

  The **Database Settings** option in the **Data Management** sub-menu of the **Configure** menu will enable the administrator to modify the database settings of the eG manager.

- **What are the database settings that can be modified?**

  The database of the eG manager stores the instantaneous measurements reported by the agents, hourly, daily and monthly trends that summarize the measurements made in the past, a history of past alarms (events) including their descriptions, their start and stop times, audit logs, the measure data used for generating reports, the detailed diagnosis measures (if applicable), and

configuration data and details of configuration changes (if the eG license enables Configuration Management). The **DATABASE PURGE PERIODS** and the **DATABASE CONNECTION POOL SETTINGS** are the settings that can be configured via the eG user interface. You can also schedule day-end activities such as trend computation, scheduled mail generation, and cleanup, by specifying the time in the **Day-end activities start at this time** field.

To optimize accesses to the database, the eG manager uses connection pooling. By using a pre-established set of connections and multiplexing requests over these connections, the eG manager ensures that individual connections are not established and closed for each request. The **DATABASE CONNECTION POOL SETTINGS** govern the initial and the maximum number of connections in the connection pool.

- **Can I configure specific cleanup frequencies for detailed diagnosis data reported by individual tests?**

  Yes. For this, navigate to the **DETAILED DIAGNOSIS DATA PURGE PERIODS** page by clicking on the **Advanced Settings** button in the **DATABASE SETTINGS** page.

- **How do I view the details of the eG database?**

  Use the menu sequence: Configure -> Data Management -> Database Properties for this purpose.

- **How can I configure the mail settings for the eG manager?**

  The mail settings can be configured via the options provided by the **Mail** sub-menu of the **Alerts** menu.

- **By default, the eG Enterprise system transmits email alerts using the SMTP protocol. Can I change this default setting?**

  Yes, you can. In the **MAIL/SMS SETTINGS** page that appears when the Alerts -> Mail -> Server menu sequence is followed, will contain a **Mail Protocol** list. **SMTP** is the default selection in this list. You can, if you so need, change this default setting by picking the **SMTP-TSL** or **SMTP-SSL** options from this list. If the mail server through which you wish to send the mail messages is **SSL-enabled**, then select, **SMTP-SSL** from the **Mail protocol** list box. If your mail server offers enhanced security and provides certificate based authentication, select the **SMTP-TLS** option from the **Mail protocol** list box.

- **Can I configure a custom subject for mails?**

  Yes. The administrator can customize the subject of the alarm mails by specifying an appropriate subject. Towards this end, the administrator will first have to indicate whether he/she needs to provide a simple, brief subject, or a more descriptive subject. To provide a short and crisp subject, select the **Concise** option against **Mail subject format**. In which case, the administrator has to specify the mail subject in the **Mail subject** text box. To provide an elaborate subject, which should include the names and/or types of components to which the alarms pertain, select the **Descriptive** option. In this case, administrators will be allowed to "build" the entire mail subject, by first specifying how the subject should begin in the **Start of mail subject** text box. Then, he/she should select the **Contents of mail subject** from the list box. The administrator can choose to include either the problem component names, or the component types, or both in the mail subject. The maximum number of problem components to feature in the mail subject should also be indicated in the **Maximum components in mail subject** text box.

- **What is the significance of the 'Alternative Mail sender ID' specification in the MAIL/SMS SETTINGS page?**

  By default, eG Enterprise sends email alerts from the **eG Administrator mail ID** configured in this page. In MSP environments typically, different support groups are created to address performance issues relating to different customers. These support groups might prefer to receive problem intimation from customer-specific mail IDs instead of the global admin mail ID, so that they can instantly identify the customer environment that is experiencing problems currently. Moreover, this way, every support group will be enabled to send status updates on reported issues directly to the concerned customer, instead of overloading the admin mailbox. To facilitate this, the **MAIL/SMS SETTINGS** page allows the administrator to configure multiple **Alternative Mail sender IDs** - normally, one each for every customer in case of an MSP environment. While configuring multiple sender IDs in the space provided, ensure that you press the **Enter** key on your keyboard after every mail ID. This way, every ID will occupy one row of the text area. Later, while creating a new user, the administrator can select one of these configured sender IDs from the **Mail sender** list in the **ADD USER** page, and assign it to the new user. This ensures that all email alerts received by the user are generated by the chosen ID only.

- **If multiple alarms are generated simultaneously, then, by default, how does the eG manager send email alerts to the configured recipients? Can this default setting be overridden?**

  If multiple alarms are generated simultaneously, then the eG Enterprise system, by default, sends a single email alert comprising of all the alarm information. Accordingly, the **SEND SEPARATE MAIL FOR EACH ALERT** flag is set to **No** by default. To ensure that a separate email is sent for every alarm, select the **Yes** option.

  Note:

  The "separate email/SMS alert" flag setting will take effect only if a user is configured to receive email/SMS alerts for the **NEW** alarms. For users who are configured to receive the **COMPLETE LIST** of alarms, details of multiple alarms will continue to be clubbed in a single email/SMS regardless of the flag setting.

- **The mail server that I have configured for the eG manager requires a user to login before sending mails. How do I configure this authentication?**

  Select the **Yes** option against the **Does SMTP server require authentication?** flag in the **MAIL SERVER SETTINGS** page. Then, provide a valid **USER** name and **PASSWORD** for logging in.

- **Can I receive email alerts when the state of a component changes to Unknown?**

  Yes. You can receive Unknown state mail alerts by configuring a setting in the **eg_services.ini**. To do so, follow the steps given below.

  1. Edit the **eg_services.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory to set the values of the parameters in the **[UNKNOWN_STATUS_REPORTING]** section, as shown below:

     ```
     UnknownStateMail=No

     UnknownStateSMS=No
     ```

```
UnknownStateInfoMail=

UnknownStateInfoSMS=

UnknownStateMailList=

UnknownStateSMSList=

DefaultUnknownStatePeriod=
```

In order to receive unknown state mails, set the **UnknownStateMail** parameter to **Yes**. The default is **No**. Similarly, specify **Yes** against **UnknownStateSMS** parameter to be able to receive SMS alerts when the state of a component changes to Unknown. The eG Enterprise system can also be configured to send out unknown state alerts even if the state of a test descriptor changes to Unknown, by setting the **UnknownStateInfoMail** parameter to **Yes**. To receive SMS alerts to that effect, set **UnknownStateInfoSMS** to **Yes**. You can even specify the users who should receive the email/SMS alerts by providing a comma-separated list of mail ids and mobile numbers (as the case may be) against the **UnknownStateMailList** and **UnknownStateSMSList** parameters, respectively. The eG Enterprise system will email/SMS the configured users only when a component remains in the Unknown state for the duration specified in the **DefaultUnknownStatePeriod** parameter. This duration can also be set specific to a test, by inserting the test name as a parameter in the **[UNKNOWN_STATUS_REPORTING]** section and providing a value against it, as shown below:

```
OraTableSpacesTest=60
```

The **DefaultUnknownStatePeriod** will automatically apply to all those tests which do not have a specific unknown state period or which have been misspelt in the **[UNKNOWN_STATUS_REPORTING]** section.

2. Finally, save the **eg_services.ini** file.

- **Can I configure the eG Enterprise system to notify me when a problem gets fixed?**

Yes. The eG Enterprise system can also be configured to send email alerts/SMS when a problem gets fixed. This can be done by setting the **Send mails/SMS when alarms are cleared** flag in the **MAIL ALERT PREFERENCES** page to **Yes**.

- **How do I make sure that the detailed diagnosis information is also sent along with the email alerts?**

To ensure that the detailed diagnosis information accompanies the alarm details, set the **Include detailed diagnosis(DD) in mail alerts** flag in the **MAIL ALERT PREFERENCES** page to **Yes**.

- **What type of manager settings can be defined in the MANAGER SETTINGS tree of the SETTINGS page in the eG administrative interface?**

The **MANAGER SETTINGS** tree enables you to define the default settings for a few critical operations of the eG manager, and also allows you to indicate the manager functions for which logging is to be enabled. To access this page, select the **Settings** option from the **Configure** menu in the eG

administrative interface and expand the **Manager Settings** node in the tree structure in the left panel of the page. For more details, refer to the *eG User Manual*.

- **What type of monitor settings can be defined in the MONITOR SETTINGS tree in the SETTINGS page of the eG administrative interface?**

  To enable you to configure the default settings for the monitor and reporter interfaces, eG Enterprise offers the **Monitor Settings** tree. To access this page, select the **Settings** option from the **Configure** menu in the eG administrative interface and expand the **Manager Settings** node in the tree structure in the left panel of the page.. For more details, refer to the *eG User Manual*.

- **How do I configure the measures that need to appear in the 'Measures At-A-Glance' section of the eG monitoring dashboard?**

  By default, the Monitor Dashboard provides a **Measures At-A-Glance** panel, which allows users to view the min/max values of critical measurements updated in real-time. Users can thus receive instant status updates on sensitive performance parameters, and can also accurately determine, at a single glance, the pain points of an infrastructure. This useful display is governed by the **Compute top metrics** flag that appears when the **Enable/Disable Metrics** sub-node of the **Measures At-A-Glance** node in the **Monitor Settings** tree is clicked. Setting this flag to **Yes** will enable the **Measures At-A-Glance** panel in the eG monitoring console. To configure the measures that are to be listed in the **Measures At-A-Glance** panel, click on the **Add a New Measure** sub-node under the **Measures At-A-Glance** node in **Monitor Settings** tree.

6.1 Detailed Diagnosis Configuration

- **What is detailed diagnosis?**

  To make diagnosis more efficient and accurate, the eG Enterprise Suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. For example, when the CPU usage of a host reaches the threshold, the agent can be configured to provide more details – e.g, the top 10 process that are consuming more CPU resources. Optionally, this capability can also be configured to periodically generate detailed measures, regardless of the occurrence of problems.

- **Can I configure the frequency with which detailed measures are generated? If so, how do I do it?**

  Yes. Using the Agents -> Settings -> Detailed Diagnosis sequence, you can define the frequency with which detailed measures are to be generated.

  In the **Diagnosis period during normal operation** text box of the **DIAGNOSIS SETTINGS** page, specify the frequency with which the eG agents need to provide a detailed diagnosis of a measure, regardless of its state (i.e. good or bad). In the **Diagnosis period during abnormal operation** text box, specify the frequency with which the eG agents should report detailed diagnosis measures if they detect a problem.

- **What happens if I set both the normal and abnormal frequencies to 0?**

  By setting both the normal and abnormal frequencies to 0, you can disable the detailed diagnosis capability for all the tests executed by the eG Enterprise system.

- **Can I enable the Detailed Diagnosis capability of a test for specific components/component-types?**

  Yes. To enable the detailed diagnosis capability for a test across all components of a type, follow the menu sequence: Agents -> Tests -> Detailed Diagnosis -> Enable / Disable. Select a **Component type** from the page that appears next. All the tests mapped to the chosen component-type for which DD is currently enabled will appear in the **DD ENABLED TESTS** list. All the tests mapped to the chosen component-type for which DD is currently disabled will appear in the **DD DISABLED TESTS** list. To enable DD for a test, pick the test from the **DD DISABLED TESTS** list and click the **>>** button; this will move the selection to the **DD ENABLED TESTS** list. Doing so will ensure that DD is enabled for the chosen test across all components of the selected type.

  To enable DD for a specific component of a type, follow the Agents -> Tests -> Configure -> Specific menu sequence. Pick the **Component type** and the **Component** for which tests are to be configured. Pick the test to be configured from the **UNCONFIGURED TESTS** list, and click the **Configure** button to configure it. The configuration parameters of the test will then appear. If DD is enabled for this test across all components of the chosen type, then. an additional **DETAILED DIAGNOSIS** flag will appear. Set this flag to **Yes** to enable DD for this test for the chosen component.

6.2 Supermanager Configuration

- **What is a Supermanager?**

  Large enterprises often have thousands of devices, servers, and applications that have to be managed, and a single eG management console may not have the capacity to handle the entire enterprise. To support such enterprises, multiple eG managers may be needed. However, if each of these managers operates independently, they may not provide a common view of the entire enterprise. Hence, it could be very cumbersome to have the IT staff of the enterprise login to different eG management consoles to get a complete view of the status of the target infrastructure.

  A SuperManager is a manager of managers that provides a consolidated view of the status of the IT infrastructure that is being handled by different eG managers.

- **What are the different types of Supermanagers that can be configured using the eG administrative interface?**

  The eG Enterprise Suite offers two options for configuring a super manager. The **eG SuperManager** is a 100% web-based component of the eG Enterprise Suite that provides a consolidated view across disparate eG managers. On the other hand, an administrator can also use the Computer Associates Network and System Management (NSM) product as a super manager (i.e., the **CA Supermanager**).

6.3 Configuring Reports

- **What is the eG Reporter?**

  The eG Enterprise system embeds an optional eG Reporter component that allows monitor users to generate a wide variety of reports. Using the eG Reporter, users can view reports that provide insights into the performance of the critical components of their infrastructure. The eG Reporter

enables analysis of collected metrics for trending, capacity planning, problem diagnosis, or service level audits.

- **The eG manager already embeds a robust graphing capability. Why do I need a separate component called the eG Reporter?**

  While the eG manager allows specific measurements to be plotted and analyzed in isolation, eG Reporter allows multiple measurements to be correlated and analyzed simultaneously.

- **How do I configure the reports?**

  Reports can be configured using the Configure -> Reports submenu.

- **What are the options offered by the Reports menu?**

  The options offered by the **Reports** menu are:

  - **Default**: The eG manager includes a number of pre-canned report templates. These templates define the set of metrics that are plotted together on the same timeline for cross-correlation and analysis. These default report templates can be refined using the **Default** option.

  - **Specific**: In many cases, an administrator may want to specify a different set of metrics for each component being managed. For example, the reports for the Oracle server running on the production server could be different from that running on the staging server. Such reports that are specific to a particular component/application/site are known as **Specific Reports**. Such reports can be configured using the **Specific** option.

  - **User**: eG Enterprise also provides administrators the ability to define reports for individual users. These user reports can be configured using the **User** option.

  - **Consolidated**: In order to enable you to focus on the critical/hyper-sensitive attributes of server/zone performance, eG Enterprise allows you the flexibility to choose the measures that are to be displayed in a **Zone**, **Server**, or **Service Consolidated Report**. In addition, the default **Time period**, **Timeline**, and **Weekend** settings for **Consolidated** and **Thin Client** reports can also be set. These settings can be defined using the **Consolidated** option.

  - **Capacity planning** : To enable users to effectively plan the capacity of their environment, eG Enterprise allows administrators to configure the trend manager to compute percentile values, and then generate detailed *Capacity Planning* reports using the computed percentiles. These capacity planning reports and the percentile values to be computed can be configured using the **Capacity Planning** sub-menu of the **Reports** menu.

For more details on report configuration refer to the *eG User Manual*.

- **What are the different types of reports that can be generated using the eG Reporter?**

  The reports generated by the eG Reporter can be broadly classified as follows:

  - **Executive reports**: These reports offer a service level summary for each component of the IT infrastructure. This summary depicts the percentage of time that a component has been problem free. The main intent of these reports is to help senior management of an organization identify which components or domains (i.e., component types) are most problem-prone in the IT infrastructure.

- **Operation reports**: While executive reports provide a high-level overview of the performance of an infrastructure component, the operation reports provide detailed insights into specific measurements and how they compare with their threshold bounds. While a number of pre-canned reports are available for each component being monitored, administrators also have the flexibility to build customized reports.
- **Comparison reports**: While operation reports allow analysis of different metrics of a specific component, comparison reports can be used for cross-correlation – i.e., to compare different metrics across different components. For example, the response time of the network, server, database, application, etc. can be compared and plotted in the same report.
- **Snapshot reports**: These reports are ideal for post-mortem analysis. If an administrator is aware that a problem occurred in the past during a certain period of time, he/she can use the Snapshot reports to quickly identify which measurements for a specific component were not in conformance with their thresholds.
- **Consolidated reports**: Use these reports to assess the performance of chosen server farms/servers/services over time. Periodic analysis of overall performance enables administrators to accurately predict behavioral trends and sizing requirements, and thus perform effective capacity planning.
- **Thin client reports:** These reports facilitate root-cause analysis and identification in Citrix/Terminal/VDI server farms. Thinclient administrators can use these reports to assess the resource utilization of specific applications in the farm, analyze user behavior across the farm, detect the open sessions in the farm, etc.
- **Event Log Compliance Report:** In order to comply with emerging standards such as HIPPA, Sarbanes Oxley, etc., IT infrastructure operators require a wide range of reports that reveal the service level acheived by a target environment, intermittent breaks in service delivery (if any), the reasons for the breaks, how quickly was the service restored, etc. To cater to these compliance requirements, the eG Reporter offers **Event Log Compliance** reports.
- **Virtualization Report:** To enable users to oversee the performance of their virtualized infrastructure over time, and to analyze the resource usage patterns of VMs and physical servers deployed on various virtualization platforms, eG Reporter now includes **Virtualization** reports.
- **Capacity Planning**: To enable users to effectively plan the capacity of their environment, eG Enterprise allows administrators to configure the trend manager to compute percentile values, and then generate detailed *Capacity Planning* reports using the computed percentiles.

➢ **How does the eG manager prioritize the different types of Operation reports?**

The eG manager assigns the highest priority to the user-specific configuration. Therefore, such a configuration will override default or host/application/site-specific configurations. If specific report configurations exist for a chosen component, then they will override the default configurations. Similarly, if no user/site/application/host-specific reports have been configured, then the default configurations will apply automatically.

6.4 Maintenance Policies

- **What are Maintenance Policies, and why do we need them?**

  During times when an IT infrastructure is under maintenance, it is but natural that few/all of the monitored components are rendered unavailable. This in turn, could cause the monitoring tool to generate a plethora of alarms indicating a "non-existent" problem situation. In order to prevent such an occurrence, eG allows administrators to:

  - Define maintenance policies (using the Configure -> Maintenance -> Define Policies menu) based on the periodicity of the maintenance procedures performed on the environment
  - Group the defined policies (using the Configure -> Maintenance -> Define Groups menu)
  - Associate the groups with components/hosts in the target environment (using the Configure -> Maintenance -> Associate Groups menu)
  - These features enable administrators to "switch off" eG alarms for specific components during maintenance periods.

- **What do I associate maintenance policies with?**

  You can associate maintenance policies with hosts, applications, or specific tests.

- **When do maintenance policy configuration changes take effect?**

  Changes made to maintenance policies are instantly propagated to the database.

- **I have configured a maintenance policy for a component in the eG admin interface. Is it possible to know about it in the eG monitoring console?**

  Yes, it is possible to know whether the component/host/test is under maintenance in the eG monitoring console. A 'spanner' icon appears alongside a monitored host/component/test in the eG monitoring console, during such time a maintenance policy is active on that host/component/test.

- **Where in the eG monitoring console can I find information related to maintenance periods?**

  If you click on the **DETAILS** button in the **Measurements** panel of the layer model page of a component, the resulting **TEST DETAILS** page displays the details of the maintenance policy associated with that component/host/test.

- **Will the maintenance policy information be available always in the TEST DETAILS page?**

  No. This information will be available only for the duration for which the maintenance policy is effective. In other words, if a maintenance policy has been configured to suppress the alerts of a component between 1 PM and 3 PM every Friday, then the **TEST DETAILS** page of that component will display the maintenance information only during 1 PM and 3 PM on Fridays.

➢ **What details related to the maintenance policy will be available in the TEST DETAILS page?**

- **Maintenance Match Criteria:** If the maintenance policy applies to the monitored host as a whole, then the host name will be displayed here. If the policy applies to an application on the host, then the component name/IP will be displayed here. Similarly, if the maintenance policy suppresses the alerts related to a particular host-level test, then the test name and host name will be displayed in the format: *TestName/HostName*. On the other hand, if the maintenance policy hides the alerts pertaining to a particular component-level test, then the test name and component-name will be displayed here in the format: *TestName/ComponentName*. If multiple maintenance policies prevail during the same period, then the match criteria will be displayed as a comma-separated list.

- **Policy group name:** The policy group into which the maintenance policy/policies has been added; multiple groups appear as a comma-separated list.

- **Policy name:** The name of the maintenance policy; if multiple maintenance policies execute during the measurement time, then the policy names will be displayed as a comma-separated list.

- **Policy time**: The time specifications associated with the maintenance policy/policies displayed against Policy name; multiple time specifications are displayed as a comma-separated list.

6.5 Discovering Components

➢ **What mechanisms does eG Enterprise use to automatically discover components in the target environment?**

eG Enterprise is capable of automatically discovering the components in the target environment. To perform this discovery, administrators can use either the central eG manager or the eG agents installed on the target hosts. Both these discovery methodologies are briefly discussed below:

- **Discovery using the eG manager**: If the eG manager is used to discover the targets for monitoring, then such a discovery will typically be based on the port number(s) on which the components are listening. In case of network devices or components that do not listen on any port, the eG manager uses SNMP for discovery. When discovery is triggered, the eG manager uses a unique port scanning technique/SNMP (as the case may be) to discover all the components/network devices that are in configured IP ranges, regardless of whether agents are installed on them or not. Typically, an agent is installed on a host only if the performance of one/more applications executing on that host interests the administrator.

  Since the eG manager discovers components unmindful of whether agents are monitoring them or not, this discovery process might end up discovering a wide variety of applications that the administrator might not even be interested in monitoring! Secondly, manager discovery is based on the assumption that the eG manager has access to all the components in the target environment. In the real world however, this might not be the case. The manager could be behind a firewall, and might hence be denied access to many/all components in the target environment; in this case therefore, with manager-based discovery, a number of components could go undiscovered.

- **Discovery using the eG agents**: As stated earlier, typically, an agent is installed on a host only when an administrator is interested in monitoring one/more applications executing on that host. In large IT environments with thousands of components, it is often difficult for administrators to manually track where agents are installed, and manage only the

corresponding hosts and applications. In such environments therefore, administrators might prefer to run a discovery procedure that automatically discovers only those components that they are "interested in monitoring" - i.e., an auto-discovery procedure that can discover only those applications which are executing on the hosts where agents are installed; this ensures that the eG administrative interface is not crowded with a wide variety of applications that an administrator might not even be interested in monitoring. This can be achieved only **if discovery is performed using the eG agents**.

The agents use a port scanning technique to discover the applications executing on the hosts on which they are installed. Since every agent performs the discovery and communicates the results to the eG manager, the location of the eG manager will not in any way impact the discovery process - i.e., even if the eG manager is behind a firewall and is not able to access the components in the target environment, the eG agent will be able to promptly detect the additions/removals in the environment everytime it rediscovers, and will be able to update the eG manager with this knowledge.

Note that **since agents are not installed on network devices, network devices are not discovered by the agent discovery process**.

- **Can the solution auto-discover inter-application dependencies as well? If so, how?**

Whether an infrastructure is virtual or physical, inter-dependencies exist between applications. For example, a web server uses a middleware application server, and an application server relies on a database server. eG Enterprises uses this inter-dependency information for root-cause diagnosis – so administrators can determine where exactly the problem lies and where the effects are. In earlier versions of eG Enterprise, inter-application dependencies had to be manually configured. Not only is this time consuming, it is also prone to human errors. In many cases, administrators may not even know all the dependencies that exist in the infrastructure. An incorrectly built segment topology impairs the eG correlation engine's ability to auto-correlate performance, thereby resulting in inaccurate root-cause diagnosis.

V 5.6 of eG Enterprise includes the capability to auto-discover inter-application dependencies. This auto-discovery reduces the time and effort involved in setting up the performance monitoring solution and also reduces the human errors that can be involved in manual specification of inter-application dependencies.

Discovery of inter-application dependencies is initiated by the agents. Using a variety of approaches (e.g., reading application configurations, checking Windows registry entries, checking TCP port connectivity, etc.), each eG agent discovers inter-dependencies between applications running on the system it is installed on and other applications in the infrastructure.

The eG manager is responsible for aggregating dependency information reported to it by different eG agents. When an administrator constructs a segment, he/she now has the option to create the segment's topology manually or by using auto-discovered information. If the administrator chooses to use the auto-discovered information, he/she can choose the starting components of the segment and from there view the auto-discovered dependencies. For each dependency, the administrator has the option to delete or reconfigure the dependency. After all the dependencies have been reviewed, the administrator can save the segment and its topology. As before, this information is then used by eG Enterprise's patented correlation engine for automatic root-cause diagnosis.

- **How do I enable agent-based discovery?**

  Follow the Agents -> Settings -> Discovery menu sequence, and set the **Enable agent discovery** flag in the page that appears to **Yes**.

- **How do I discover components using the eG manager?**

  From the **Infrastructure** menu of the eG administrative interface, select **Discover**. The **START DISCOVERY** page will then appear. This page displays two panels - a left panel comprising of a **DISCOVERY** tree structure, which consists of nodes and sub-nodes that enable you to quickly navigate the discovery-related options, and a context-sensitive right panel that changes according to the node chosen from the tree. Since the **Start Discovery** node is by default chosen from the tree, the **START DISCOVERY** page is available in the right panel depicting the various inputs that need to be specified before discovery can be started. Refer to Chapter 3 of the *eG User Manual* for the details to be specified in this page. Finally, click on the **Start Discovery** button at the bottom of this page to start the discovery process.

- **How do I change the default port preferences mapped to a component?**

  Click on the **Ports** sub-node under the **Settings** node in the tree-structure in the **START DISCOVERY** page. Using the **CHANGE PREFERENCE** page that displays in the right panel, you can edit the port settings of applications.

- **Do I have to run the discovery process every time?**

  The administrator need not run the discovery process every time. Specify the **Re-discovery period** in the **START DISCOVERY** page that appears when the **Start Discovery** sub-node of the **Actions** node in the **DISCOVERY** tree structure is clicked. The **Rediscovery period** determines the frequency with which the discovery process executes. This frequency governs how quickly the eG manager is able to discover servers that may have been newly added to the target environment.

- **How do I know which components were recently discovered by the eG manager?**

  Typically, in large infrastructures, discovery/re-discovery may result in automatically discovering numerous components of varied types. Though the discovered components can be viewed in the **MANAGE / UNMANAGE** page upon the selection of a component-type, administrators of such environments would prefer an interface that can provide them with a quick look at what components have been discovered across component-types, so that they can swiftly decide which components need to be monitored and which do not require monitoring.

  To allow administrators easy and instant access to the details of newly discovered components, eG Enterprise provides the option to configure a 'discovery pop-up'. Soon after discovery, this pop-up, once enabled, will display a type-wise list of recently discovered components, thus saving administrators the trouble of navigating to the **MANAGE/UNMANAGE** page to achieve this.

- **How do I enable the discovery pop-up?**

  By default, this pop-up is disabled. To enable the pop-up, open the expand the **Manager Settings** tree in the **SETTINGS** page (Configure -> Settings menu sequence), select the **Discovered Components** node in the tree, and set the **Show pop-up** flag in the **DISCOVERED COMPONENTS POP-UP** section to **Yes**.

- **How does the eG manager discover the IT infrastructure elements?**

  The eG manager uses port scanning and SNMP for discovering the IT infrastructure elements.

- **Can I know what components of which type have been managed across the infrastructure?**

  Yes. While monitoring large environments using eG Enterprise, administrators may want to instantly figure out the total number of managed components in the environment across categories / component-types to assess the load on the eG manager and the database. Some other times, they may require a quick summary of the number of managed components within a particular zone/service/segment. To enable such administrators and top-level executives to receive an overview of the managed infrastructure based on chosen criteria, the eG administrative interface provides the **MANAGED INFRASTRUCTURE** page.

  This page, which appears when the Infrastructure -> View menu sequence is followed, provides administrators with a quick look at the managed infrastructure as a whole, based on a chosen service/zone/segment.

- **What are the other discovery options available in the DISCOVERY page?**

  You can auto-discover the vSphere/ESX servers in the environment by connecting to the vCenter server managing them. Using the **DISCOVERY** page, you can configure the vCenter server to be used for performing such a discovery. Similarly, you can auto-discover the IBM pSeries servers hosting AIX LPARs by connecting to the HMC server in the environment. To facilitate this discovery, use the **DISCOVERY** page to configure the HMC server of interest. Likewise, you can use the **DISCOVERY** page to configure an RHEV manager using which the RHEV hypervisors in the environment are to be auto-discovered. In the same manner, you can configure the *AWS EC2 Cloud* with the *AWS EC2 Regions* to be monitored.

6.6 Managing/Unmanaging Components

- **How do I manage/unmanage the discovered applications/servers?**

  Select the **Servers** option from the **Manage/Unmanage** sub-menu of the **Infrastructure** menu helps the administrator to manage or unmanage the discovered components. Refer to Chapter 3 of the *eG User Manual* for the steps involved in this process.

- **Can I manage/unmanage all applications executing on a particular host at one shot?**

  Yes. You can manage/unmanage all applications executing on a particular host using the **SYSTEMS – MANAGE/UNMANAGE** page that appears when the **System** option is chosen from the **Manage** sub-menu of the **Infrastructure** menu.

- **Can I delete a server/application that is being managed by the eG manager?**

  Yes. A server/application being managed by the eG manager can be deleted. The administrator has to choose the component to be deleted from the **MANAGED COMPONENTS** list box in the **COMPONENTS - MANAGE/UNMANAGE** module, and then click the **Delete** button at the bottom of the list box.

- **Can I delete an unmanaged component?**

  Yes. An unmanaged component can be deleted in the same way as a managed one. The administrator has to choose the component to be deleted from the **UNMANAGED COMPONENTS** list box in the **COMPONENTS - MANAGE/UNMANAGE** module, and then click the **Delete** button at the bottom of the list box.

- **What does an asterix (*) beside a component name in the Manage / Unmanage page imply?**

  When the eG manager discovers a component, it will mark the component as being unmanaged and prefix it with an asterix (*) to indicate that this component has been newly discovered. The administrator then has to explicitly choose to manage the component.

6.7 Adding Components

- **How do I add a new component explicitly to the eG Enterprise system?**

  In cases in which the eG manager's discovery process is not able to discover a specific component or set of components (like network node, load balancers etc.), the **Add/Modify Servers** option of the **Infrastructure** menu permits the user to explicitly add the component for monitoring by the eG agent. Upon choosing the **Add/Modify Servers** option, the administrator sees a display of all the types of components currently being managed. For each type of component, by choosing the **Add New Component** option, the administrator can choose to add a new component for monitoring by the eG agent. For the details regarding the parameters that have to be specified by an administrator when adding a new component for monitoring by eG agent, refer to Chapter 3 of the *eG User Manual*.

- **Should I restart the eG manager when new components are added to the target environment?**

  The eG manager is able to detect new components in the environment and is capable of adding new tests for these components without having to shut the manager down and restart it. Changes to thresholds, any parameter changes to the tests, the enabling/disabling of all tests, etc., are handled by the eG Enterprise Suite without requiring the eG manager to be started up and shut down.

- **Can I add an Oracle server just like the other types of components?**

  An Oracle server listening on a port may comprise of multiple instances. The eG manager can monitor the different database server instances individually. The database server instances associated with an Oracle server have to be configured manually via the **Add/Modify Servers** option of the **Infrastructure** menu. If the instances corresponding to the Oracle server are not configured, then a **WARNING** icon appears at the top left corner of the **MANAGE/UNMANAGE, ADD NEW COMPONENTS** and **TOPOLOGY CONFIGURATION** pages. By moving the mouse over this icon, the user can see the warning message that displays the components for which the instances have not been configured. Please look into the *eG User Manual* for how to specify the information about the Oracle server under consideration.

- **I have already specified the details of the component in the ADD MODIFY COMPONENTS page. Now if I wish to change the IP address of the component, can I do it?**

  Yes, it is possible.

Click on the **Change IP** button in **Modify Component Details** page. This way, you can modify the IP address.

- **What is Agentless Monitoring?**

  The agent-based approach to monitoring requires one internal agent per system that is to be monitored. As operating systems and applications have evolved, they have incorporated newer instrumentation mechanisms that allow monitoring of these environments from remote locations – i.e., external to the system or application that is to be monitored. The main advantage of this remote monitoring approach is that it does not require an agent to be installed on every system that is to be monitored – hence, the name "*agentless monitoring*" for this approach.

- **What are remote agents?**

  Agentless monitoring is implemented in the eG Enterprise Suite using **Remote Agents**. A remote agent is an agent that is capable of monitoring a number of systems and applications remotely, in an agentless manner.

- **How do remote agents monitor a component?**

  For monitoring Microsoft Windows systems and applications, a remote agent uses Netbios/perfmon to communicate with the operating system/applications. For monitoring Unix systems, secure shell (SSH) is used. In addition, for specific applications, the remote agent uses application-specific protocols to communicate with the application (e.g., SQLNet for Oracle databases, HTTP for WebLogic and WebSphere application servers, JDBC for Sybase, etc.). When deploying a remote agent, it is essential to ensure that the remote agent can communicate using multiple of these agentless mechanisms with the target environment (appropriate firewall rules must be configured for this purpose).

- **When can an eG user choose between the agent-based and agentless monitoring approaches?**

  Unlike many existing solutions, the eG Enterprise Suite does not require that IT administrators choose between one of these contrasting approaches at the time of installing/deploying the eG Enterprise Suite. Instead, when adding any new application/system for monitoring, administrators have a choice of whether to use agent-based or agentless monitoring for the application or system under consideration. For instance, an administrator can choose to monitor the most critical servers in an agent-based manner, and to monitor the less critical servers in the staging/development environment in an agentless manner.

- **Under what circumstances will the option to enable agentless monitoring for a specific component be available?**

  - The eG license should indicate agentless monitoring support
  - Atleast one remote agent has to be configured in the environment

- **What are the capabilities which are not available for servers or applications monitored by remote agents?**

  - Detailed diagnosis
  - Automatic corrective script execution
  - Remote control actions

- The eG web adapter for in-depth web server monitoring on Unix and Microsoft environments – i.e., web transaction monitoring for web servers is not available with agentless monitoring.

➢ **What are 'Aggregate Components'? How do I add them?**

eG Enterprise typically monitors every component of a type, separately. However, sometimes, administrators might want to receive an aggregated view of the performance of two/more components of a type. For instance, Citrix administrators might want to know the total number of users who are currently logged into all the Citrix servers in a farm, so that sudden spikes in the load on the farm (as a whole) can be accurately detected. Similarly, Windows administrators might want to figure out the average CPU usage across all the Windows servers in an environment, so that they can better plan the capacity of their Windows load-balancing clusters.

To provide such a consolidated view, eG Enterprise embeds a license-controlled **Metric Aggregation** capability. This feature, when enabled, allows administrators to group one or more components of a particular type and monitor the group as a single logical component, broadly termed as an *aggregate component*. The eG Enterprise system then automatically aggregates the metrics reported by the components in the group by applying pre-configured aggregate functions on them, and reports these metrics as if they were extracted from the managed *aggregate component*. Separate thresholds need to be set for the aggregated metrics to track deviations in the consolidated performance. The state of the *aggregate component* is governed by these exclusive thresholds, and not by the state of the components within the group.

Moreover, since only remote agents can perform metric aggregation, one/more **premium monitor licenses** would be required for implementing this capability.

Using this **Metric Aggregation** capability, administrators can perform the following:

- Effectively assess the collective performance of a group of components of a particular type
- Easily study load and usage trends of server farms (or groups) as a whole
- Accurately detect resource inadequacies or unusual load conditions in the component group or farm
- Compare and correlate the performance of the member components with that of the *aggregate component*, so that the reasons for performance issues with the *aggregate component* can be precisely determined;

The very first time the administrator manages/adds a component of a type, the eG Enterprise system dynamically creates a corresponding *aggregate component type*. For instance, if an *IIS web* server component is added to the eG Enterprise system for the first time, a component type named, *IIS Web Aggregate* is automatically created alongside. To add a component of the dynamic *aggregate component type*, use the *Infrastructure -> Aggregates -> Add/Modify* menu sequence.

6.8 Internal Agent Assignment

➢ **What is the purpose of the agent per system capability of the eG Enterprise Suite?**

By default, if a host has multiple IP addresses, the eG Enterprise system requires one agent license for each IP address that is managed internally. Likewise, if multiple nicknames are used for the same IP address, a separate internal agent license is used for each unique nickname that has

been specified. In many large environments, a single host has many IP addresses, each with different nicknames. The agent per system capability is intended to optimize the internal agent license
usage in such large infrastructures. For example, suppose a host A has two IP addresses 192.168.10.7 and 10.10.10.1, and that the first IP address 192.168.10.7 has already been managed in the eG Enterprise system. When adding the second IP address, 10.10.10.1, the administrator has the option of overriding eG's default internal agent licensing policy – in this example, the administrator can indicate that the internal agent for the IP address 10.10.10.1 is actually the one that is already associated with the IP address 192.168.10.7. By doing so, the administrator can ensure that a single agent license is sufficient to manage all the IP addresses and applications executing on a host.

- **How do I assign the same internal agent to all the IP addresses and applications executing on a host?**

  While adding a new component using the **NEW COMPONENT DETAILS** page of the eG admin interface, an **Internal agent assignment** field appears. By default, the **AUTO** option against the field will be selected. This indicates that by default, eG maps every configured IP/nick name with a separate internal agent. To manually define the IP-internal agent association, select the **MANUAL** option. Upon choosing **MANUAL**, an additional **Internal agent** list box will appear. From this list box, select the internal agent that needs to be associated with the IP of the component being created. This way, the same internal agent can be associated with multiple IPs on the same host.

- **When does the INTERNAL AGENT ASSIGNMENT field appear in the NEW COMPONENT DETAILS page?**

  The **INTERNAL AGENT ASSIGNMENT** field will appear only if the following conditions are fulfilled:

  - The eG license enables the **AGENT PER SYSTEM** flag
  - The **AGENTLESS** flag in the **NEW COMPONENT DETAILS** page is set to **NO**. In other words, users will not have the option of mapping an IP to an internal agent, if agentless monitoring is enabled

6.9 Configuring Zones

- **What is a zone?**

  Large infrastructures spanning geographies can pose quite a monitoring challenge owing to the number of components involved and their wide distribution. Administrators of such infrastructures might therefore prefer to monitor the infrastructure by viewing it as smaller, more manageable business units. In eG parlance, these small business units are termed **ZONES**. A zone can typically comprise of individual components, segments, services, and/or other zones that require monitoring. For example, in the case of an infrastructure that is spread across the UK, USA, and Singapore, a zone named *USA* can be created consisting of all the components, segments, and services that are operating in the US branch alone. The *USA* zone can be further comprised of an East-coast zone and a West-coast zone to represent infrastructure and services being supported on the two coasts of the US.

  While a service/segment contains a group of inter-related components with inter-dependencies between them, a zone contains a group of components, services, segments, or zones that may/may not have inter-dependencies.

- **How do I access the zone configuration page?**

Use the menu sequence **Infrastructure -> Zones** to access the zone configuration page.

- **What is the implication of enabling/disabling "automatic association"?**

If automatic association is enabled:

When a segment is added to a new zone, then the following (if available) will automatically be added to the zone:

- All the components within that segment
- All the services using that segment
- All the other segments that are part of the associated services

Similarly, when a segment is disassociated from a zone, then the afore-said will also be automatically removed from the zone

If automatic association is disabled:

- When a segment is added to a new zone, then only the segment and the components within the segment are added to the zone
- Similarly, when a segment is disassociated from a zone, that segment and the underlying components alone will be removed from the zone

- **Can I configure zone locations?**

Typically, zones can be used to represent the status of the IT infrastructure in a specific geographic location. eG Enterprise allows you to drill down on a geographic map to visually figure out the exact geographic area where a zone operates, and instantly evaluate the performance of the different zones spread across the different locations worldwide. To enable such an analysis, you first need to indicate the geographic location of the configured zones using the built-in map interface of eG Enterprise. For instance, while configuring an *east_coast* zone, you can use the map to point to the east coast of USA, so as to enable quick and accurate identification of zone locations. To achieve this, click on the **Add geographic location** link in the right corner of the **ZONE SELECTED FOR CONFIGURATION** section in the **CONFIGURE ZONE** page. The page that then appears reveals a world map. The geographic map display for zones is achieved through an integration of the eG Enterprise management console with the Google maps service. *You will therefore have to enable this integration in order to configure and visualize zones.* Please refer to the *eG User Manual* for steps on how to enable the integration.

6.10 Configuring Groups

- **What is a Component Group?**

Large enterprises often have load balanced groups of servers providing web, middleware, and database functionality. The servers in the group often perform the same functions and have the same set of dependencies on other infrastructure components. For a service with tens of servers in a group, the service topology representation can quickly get very cumbersome. To handle such environments, eG Enterprise allows administrators to define server groups that represent a

collection of similar servers. By including a server group in a service topology representation, administrators can indicate that all the servers in the group have the same set of dependencies on other parts of the infrastructure. Topology representation using server groups are compact and concise to represent, and simple to understand.

6.11 Configuring the Topology

- **What is the topology and why is it needed?**

The eG manager's patented correlation technology is dependent on the specification of topology information that indicates how components are interconnected and which components rely on others for their functioning. The interconnections can represent either physical connections (e.g., a web server connected to a network router) or logical dependencies (e.g., a web server using a web application server). Each interconnection is associated with a direction. The direction signifies cause-effect relationships (if any) between the components being connected together.

- **Can eG auto-discover topology? If so, how?**

Whether an infrastructure is virtual or physical, inter-dependencies exist between applications. For example, a web server uses a middleware application server, and an application server relies on a database server. eG Enterprises uses this inter-dependency information for root-cause diagnosis – so administrators can determine where exactly the problem lies and where the effects are.

eG Enterprise includes the capability to auto-discover inter-application dependencies. This auto-discovery reduces the time and effort involved in setting up the performance monitoring solution and also reduces the human errors that can be involved in manual specification of inter-application dependencies.

In the eG Enterprise system, segment/service **topology definitions** embed the **inter-application dependencies**. Discovery of this topology information is initiated by the agents.

By default, the ability of the eG agent to **automatically discover topology** is enabled. However, this default setting will take effect only if the **eG agent has the ablity to automatically discover components**. This is because, topology discovery cannot be performed without component discovery. Therefore, as soon as the **Enable agent discovery** flag in the **AGENT DISCOVERY SETTINGS** page (Agents -> Settings -> Discovery) is set to **Yes**, a **TOPOLOGY DISCOVERY SETTINGS** section will automatically appear therein, using which you can define the settings for automatically discovering the component topology using the eG agent.

The **Enable topology discovery?** flag is set to **Yes** by default, indicating that the eG agents are bundled with the ability to perform automatic topology discovery.

Once automatic topology discovery is enabled, the eG agent will run the **netstat** command on the target host every 45 minutes (by default) to determine which applications are operating on the host and which port they listen to. This default frequency can be overridden using the **Delay between successive dependency discovery attempts (Mins)** parameter in the **TOPOLOGY DISCOVERY SETTINGS** section. By default, this parameter is set to **45** by default, indicating that the default discovery frequency is 45 minutes. To override this default frequency, specify a different duration (in minutes) against the **Delay between successive dependency discovery attempts (Mins)** parameter.

Whenever the eG agent on a host runs **netstat**, it retrieves a list of ports that are operating on that host. While some of these TCP ports may be standard listening ports - i.e., TCP ports at which the applications executing on that host listen for requests from remote hosts/applications - a few other TCP ports may be local ports created dynamically on the host for a temporary purpose. To clearly differentiate between listening ports and local ports, the eG agent does the following:

- By default, the eG agent compares the ouput of four consecutive executions of the **netstat** command on a host. If a port number is repeated in all the four **netstat** outputs by default, then that port number is counted as a 'server listening port'. This default behavior is governed by the **Dependencies must be present** and the **Number of dependency discovery attempts** parameters in the **TOPOLOGY DISCOVERY SETTINGS** section. Both these parameters are set to **4** by default. The **4** against the **Number of dependency discovery attempts** parameter indicates the maximum number of consecutive **netstat** outputs to be considered for identifying the server listening ports. The **4** against the **Dependencies must exist** parameter indicates the minimum number of **netstat** outputs in which a port number should appear for it to be considered as a 'server listening port'. You can change this if you need. For instance, if you set the **Dependencies must exist** parameter to 3 and let the **Number of dependency discovery attempts** parameter to remain at **4**, then the eG agent will count the port numbers that appear in at least **3** out of **4** consecutive **netstat** outputs as active listening ports.
- Once the listening ports are identified, the agent then closely observes traffic to and from a 'server listening port', identifies the remote applications that frequently connect via this port, and thus automatically discovers the inter-relationships that exist between applications in an IT infrastructure. The interdependencies that are so discovered are then sent to the eG manager. On the other hand, all those port numbers that do not conform to the specification governed by the **Dependencies must be present** and the **Number of dependency discovery attempts** parameters are counted as local ports. All traffic to local ports are hence disregarded for the purpose of topology auto-discovery.

By default, the whole cycle of operations - beginning with isolating the listening ports to discovering inter-dependencies to reporting the discovery to the eG manager - takes 180 minutes (as indicated by the default value **4** against the **Dependencies must be present** setting) to complete. After which, the eG agent will wait for another 360 minutes (i.e., 6 hours, by default) to rediscover the topology. In other words, all the aforesaid activities will be performed again by the eG agent 6 hours after the first cycle is complete. Like other default settings, you can also override the frequency with which topology is rediscovered by the eG agent. For this, use the **Topology rediscovery period (Mins)** parameter, which is set to **360** by default.

- **Can all eG agents perform topology discovery?**

  eG Agents deployed on the following operating systems can alone perform topology discovery:

  - Linux
  - Windows (2000/2003/2008/Vista/XP/7)
  - AIX

  Solaris and HPUX agents cannot auto-discover the topology.

- **Differentiate between the USES and CONNECTS relationship.**

  The **CONNECTS** dependency indicates flow of data (e.g., a physical connection between a web server and a network router), and the **USES** dependency refers to a logical dependency (e.g., a web server and a web application server, a web application server and a database server, etc.). The key difference between the two forms of dependencies is that when one component **USES** another, problems with the latter component can actually reflect in problems with the former component. With the **CONNECTS** dependency, there are no such cause-effect relationships between the two components.

- **Where do I get to see the topology preview?**

  While the administrator is defining the dependencies within a target environment using the **CONFIGURE TOPOLOGY** page, the **SEGMENT PREVIEW** section at the end of the page displays every step of the topology specification as it proceeds. The feedback provided by the topology preview can be used by the administrator to refine the topology specification.

- **Does the topology indicate network connectivity?**

  The topology indicates how application components rely on other components for their functioning. The interconnections can represent either physical connections (e.g., a web server connected to a network router) or logical dependencies (e.g., a web server using a web application server).

- **I want to manually construct the segment topology? How do I know the application relationships to construct a topology?**

  There are various ways in which application dependencies may be maintained. Some applications may store their info in application specific log files. Others may use the system registry to store this info. Some others may maintain this info in application-specific registries. If you have enabled the eG agent's ability to auto-discover the topology, then the agent, using its built-in intelligence, accurately captures how two applications interact with each other; this, saves you the time involved in manually understanding the inter-relationships between the applications in an environment hare and in translating this understanding into the topology.

- **Can I associate a segment with a zone? If so, how?**

  Yes, a segment can be associated with a zone. This can be performed in 2 ways:

  While configuring a zone, you can add specific segments to the zone

While configuring a segment, you can associate the segment with a zone

- **While configuring 'seg_a', I associated it with 'my_zone'. What are the implications of this association?**

  If, during segment configuration, you associate the segment with a zone, then such a segment can be constructed using only the components that are already a part of the zone.

- **During segment configuration, I realized that a component that needs to go into the segment is not available in the associated zone. What do I do?**

  Any segment that is associated with a zone or service, will carry an **ASSOCIATION** section in its configuration page, which not only lists the zones and services that are mapped to the segment, but also 'links' to the corresponding zone and service configuration pages. Accordingly, the **ASSOCIATION** section in your case should indicate the zone(s) associated with your segment. Clicking on the **Zone** in question will enable you to reconfigure the zone and add the missing component to the zone.

- **Can a component be added to more than one segment?**

  By default, the components in a segment will not be candidates for inclusion in another segment. However, you can share components across segments by selecting the **Use components from other segments** check box in the topology configuration page.

6.12 Configuring Services

- **Can the eG Enterprise Suite monitor specific web sites?**

  Rather than monitoring individual elements, the eG Enterprise Suite monitors the services that end users see. Users do not access web servers. Rather, they access **web sites** that are hosted on one or more web servers. A web site offers one or more services to its users. The various services that users can avail via a web site are referred to as **transactions**. Web sites and web servers have a many to many relationship among them. In an ISP hosting scenario, many web sites may be hosted on the same web server. Alternatively, very large eBusiness sites may use multiple web servers to support a single web site. The eG Enterprise Suite supports both of these scenarios.

- **How do I configure a new site for monitoring?**

  To configure a web site, an administrator has to choose the **Topology** option from the **Services** sub-menu of the **Infrastructure** menu in the eG administrative interface. Using the **Add New Service** option in the page that appears, new web sites can be added for monitoring by eG. In the screen that appears next, the name of the new service has to be specified. Additionally, users need to indicate whether the service is a web site or not, by selecting the **Yes** option from the list box in the page. For in-depth details about configuring the dependencies of a website via the user interface, please refer to Chapter 3 of the *eG User Manual*.

- **What are the restrictions for a site name?**

  While specifying the name of a website, care should be taken to see that

  1. it does not contain any special characters that are not permissible for a site name

2. a "/" does not occur
3. does not have letters in upper case

- **Can the eG Enterprise Suite monitor services other than web sites?**

  Yes. The eG Enterprise Suite can monitor services other than web sites. To configure a service, an administrator has to choose the **Topology** option from the **Services** sub-menu of the **Infrastructure** menu in the eG administrative interface. Using the **Add New Service** option in the page that appears, new services can be added for monitoring by the eG Enterprise system. The name of the new service has to be specified in the page that appears next. Additionally, an administrator will have to indicate that the service being configured is "not" a web site by selecting **No** from the list box therein. For in-depth details about configuring the dependencies of a service via the user interface, please refer to Chapter 3 of the *eG User Manual*.

- **Can I associate a service with a zone? If so, how?**

  Yes, a service can be associated with a zone. This can be performed in 2 ways:

  While configuring a zone, you can add specific services to the zone

  While configuring a service, you can associate the service with a zone

- **While configuring 'service1', I associated it with 'my_zone'. What are the implications of this association?**

  If, during service configuration, you associate the service with a zone, then such a service can be constructed using only the components/segments that are already a part of the zone.

- **During service configuration, I realized that a segment that I want to associate with the service is not available in the associated zone. What do I do?**

  Any service that is associated with a zone or service, will carry an **ASSOCIATION** section in its configuration page, which not only lists the zones and segments that are mapped to the segment, but also 'links' to the corresponding zone and segment configuration pages. Accordingly, the **ASSOCIATION** section in your case should indicate the zone(s) associated with your service. Clicking on the **Zone** in question will enable you to reconfigure the zone and add the missing component to the zone.

- **Can I group services?**

  Large organizations may have multiple services grouped under different business units. There may hence be a need to represent groups of services as an entity. To address this requirement, eG Enterprise allows the configuration of service groups in the eG admin interface, and represents the real-time state of the service groups in the eG monitoring interface. To configure a service group, follow the menu sequence, Infrastructure -> Services -> Groups in the eG administrative interface.

6.13 Configuring Transactions

- **How do I add new transactions?**

  For each web site that has been configured, the eG agent has the ability to monitor individual transactions that happen via the web site. By choosing the **Transactions** menu option under **SERVICES** configuration menu, an administrator can view the transactions that have been currently configured for each web site. Clicking on the **ADD/DELETE TRANSACTION** button will enable an administrator to add new transactions for a web site. Please go through Chapter 3 of the *eG User Manual* for specifying the details of transactions to be added.

- **Can I delete existing transactions?**

  The **Transactions** option under **Services** sub-menu of the **Infrastructure** menu will enable the administrator to view the transactions that have been configured. From here, click the **ADD/DELETE TRANSACTION** button beside the transaction to be deleted. Finally, click the **DELETE** button to remove the transaction permanently.

- **Based on what criteria do I define transactions?**

  There are two criteria that an administrator can use to define transactions that must be monitored:

  a. Administrators can configure transactions to reflect the key operations performed by users of the web site like login, registration, browsing of the product catalog, searching the catalog, adding to shopping cart, deleting items from the cart, payment, etc. By monitoring individual transactions, web site operators can determine patterns of user accesses to individual transactions. Moreover, errors and response time issues with individual transactions can be monitored.

  b. Transactions can also be configured so as to differentiate between requests to the front-end web server and requests to the backend. For example, considering a web site that uses the iPlanet application server, all requests to the back-end application server can be represented by the pattern `*/cgi-bin/gx.cgi*` where * indicates zero or more characters. Using this approach, a web site operator can track requests sent to the back-end independent of requests targeted at the front-end web server and detect problems associated with the back-end easily.

6.14 Configuring Tests

- **How do I configure the tests for different component types?**

  For each component type, the eG Enterprise Suite has a set of predefined tests. These tests can be configured via the **Configure** sub-menu of the Agents -> Tests menu. The **Configure** menu provides two options, namely - **Specific** and **Default**. While many of the eG tests do not require any manual configuration, some tests require explicit, manual configuration. To configure such tests

for all components of a type, pick the **Default** option. To configure a test for a specific component of a type, pick the **Specific** option.

Regardless of the option chosen from the **Configure** menu, a page displaying the list of tests that are not configured appears. Choosing any of the unconfigured tests will lead the user to a screen from where the user can proceed to configure the test.

- **How do I configure the tests for an Oracle database server?**

As already mentioned, eG agents can collect a variety of statistics regarding Oracle database servers. While configuring the Oracle tests, a link is provided along with the parameters to be configured. The user is taken to the **ADD NEW USER** page by clicking on the link **Click here** above the list of parameters to be configured. The host name, port number and the sid of the Oracle server to be configured appear in the corresponding location. To add a new user, specify the database administrator's login and password. While creating a new user, the account has to be associated with a default tablespace where the user's data is hosted and a temporary tablespace which is used for buffering, sorting etc. The identities of the default and temporary tablespaces have to be input in this page. Please look into Chapter 3 of the *eG User Manual*.

- **Should I restart the eG manager every time a user changes the configuration of a test or a threshold?**

The manager or the agents do not need to be restarted every time a test configuration or thresholds for the measurements have been changed. The eG agent and manager automatically discover changes that have been made by users to the test configuration and/or measurement thresholds.

- **Can I add 'aggregate tests' for aggregate components? If so, how?**

Yes. For each aggregate component-type you manage, you can add aggregate tests, which automatically execute configured functions (average, sum, sum of average, max, min, etc.) on chosen measures to report meaningful aggregated metrics of performance. To add an aggregate test, follow the Infrastructure -> Aggregates -> Add/Modify Tests menu sequence. To configure these tests, follow the Agents -> Tests -> Configure -> Specific or Default menu sequence, just like in the case of non-aggregate tests. For more details, refer to the ***eG User Manual***.

6.15 Configuring Thresholds

- **How do I configure thresholds for various measurements?**

Thresholds can be configured via the **Default/Specific** option in the **Thresholds** sub-menu of the **Alerts** configuration menu. This will take the user to page that displays the tests that map to the selected component type. The **Default Thresholds** option beside a test enables the administrator to configure the default threshold values for the measurements pertaining to that test for the selected server type. The **Specific Thresholds** option on the other hand allows the user to configure the thresholds for individual servers. The **Refine** button in the **Specific Thresholds** page allows users to define descriptor-specific thresholds for individual server. For greater details, refer Chapter 3 of the *eG User Manual*.

- **What are the thresholding policies that the eG Enterprise Suite supports?**

The eG Enterprise Suite supports three threshold policies:

a. **Absolute thresholding:** The eG Enterprise Suite permits operators to specify time-invariant lower and upper bounds for each measurement. The values specified by the user depending upon their infrastructure in the **Minimum** and **Maximum** columns form the bounds. The state of a measurement is abnormal if the instantaneous values of the measurement cross the lower and upper thresholds for a sustained period of time (determined based on the window size and the number of threshold crossings).

b. **Relative thresholding:** eG agents track past history for each measurement and using statistical quality control techniques, automatically determine the lower and upper bounds for each measurement. Since the values of many measurements vary from time to time, the historical thresholds are computed in a time-varying manner (e.g., the normal rate of connections to an Intranet web server may be high during office hours, and low during non-office hours). As in the case of absolute thresholding, the state of a measurement is determined by comparing instantaneous values with historical thresholds. Currently, the eG Enterprise Suite supports only one policy for relative thresholding – sqc standing for statistical quality control. Future versions may support additional policies.

c. **None:** If the threshold policy for a measurement is none, an eG agent will stop tracking the state of this measurement (i.e. The agent will continue to collect values for this measurement but will not generate any alarms relating to this measurement) Even in the case of absolute and relative thresholding, the eG Enterprise Suite allows either the minimum or maximum threshold values to be "none".

- **What are multiple thresholds and how do I define them?**

The eG Enterprise system offers three levels of thresholds that correspond to the three alarm priorities - **Critical, Major, and Minor**. The user has to specify three maximum and/or three minimum threshold values in the format: *Critical/Major/Minor*. While the maximum thresholds are to be provided in the descending order, minimum thresholds have to be specified in the ascending order. For example, take the case of the Pct_disk_util measure of a host. This measure reports the percentage of disk space that has been utilized. The user can set a single Maximum threshold of say, 98, and expect to be alerted when the disk utilization crosses 98%. Alternatively, the user can also set multiple maximum thresholds, thereby instructing the eG Enterprise system to send different types of alerts at various levels of disk usage - in other words, the user can instruct the eG Enterprise system to trigger a **Minor** alert if the disk utilization crosses 50%, a **Major** alert if the disk utilization crosses 75%, and a **Critical** alert if it falls beyond 90%. This provides ample opportunity to the user to identify and attack problems early in its life cycle.

In the case of the Pct_disk_util measure in our example, the Maximum thresholds can be defined as *"90/75/50"*. Since an absolute minimum threshold is not required for the Pct_disk_util measure, it can remain as *"none"*.

According to this specification, if the Maximum threshold of 90 is violated, then a **Critical** priority alarm will be generated. This is indicative of a critical issue with the host. Similarly, if the value of this measure crosses the Maximum threshold of 75, then a **Major** priority alarm will be generated.

This is indicative of the existence of a major issue with the host. Likewise, a value beyond the Maximum threshold of 50 will result in a **Minor** priority alarm.

Note:

To omit a particular alarm priority/priorities from a multiple threshold specification, replace the value that corresponds to that priority with a hyphen (-). For example, to omit the **Major** priority alarm for the *Disk utilization* measure, the threshold specification should be: "90/-/50".

- **How do I define thresholds for a measure?**

Typically, depending upon the nature of the performance metrics reported by the eG agent, administrators may set either **Absolute** or **Relative** thresholds for the metrics. For instance, while **Absolute** thresholds can be set for time-invariant measures such as disk space usage, dynamic parameters such as connections to a web site are better governed by relative/automatic thresholds. A key advantage of automatic thresholds is that the administrators do not have to spend endless hours configuring thresholds for each and every metric. While auto-thresholding is ideal for infrastructures with periodic workloads, this can result in false alerts when used in infrastructures where the workload is not as periodic. To overcome this limitation, eG Enterprise allows administrators to set a combination of fixed and automatic thresholds for each metric. For example, take the case of the number of connections to a target server in an IT infrastructure. Since this measure increases during the working hours and decreases during the non-working hours of the infrastructure, the maximum threshold for this measure can be determined only using 'sqc' - i.e., the relative thresholding policy. Accordingly, a relative thresholding policy was applied to this measure. Now, say that there was no user activity in the target infrastructure for over 15 days owing to the Christmas holidays. When the users returned, connections to the target server, naturally, peaked! However, since eG Enterprise auto-computed the maximum threshold for this measure based on the performance data of the last 15 days (when there was nearly zilch user activity), it treated the sudden increase in activity as a violation, and generated a plethora of inconsequential alarms! To avoid such false alerts, you can configure an **Absolute Maximum** and a **Relative Maximum** threshold for the number of connections measure. A fixed threshold applied along with an automatic threshold specifies a realistic bound that has to be crossed before an alert is to be triggered. eG Enterprise then compares the actual value with the higher of the two thresholds, and will generate an alert only when the higher threshold is violated. For instance, you can set an **Absolute Maximum** threshold of 100 for the number of connections measure, and a **Relative Maximum** threshold of 'sqc'. In such a situation, when, after a brief period of inactivity, operations resume in a target infrastructure, eG Enterprise first compares the 'near-zero' relative threshold with the absolute threshold of '100', and fixes the higher value of 100 as the threshold for the number of connections measure. It then compares the current level of activity on the target server with the threshold of 100, thereby generating alarms only when the number of connections are in excess of 100. Similarly, if an **Absolute Minimum** and a **Relative Minimum** threshold is set for any measure, then eG Enterprise will generate alarms only when the current value falls below the lower of the two threshold settings.

- **Is it possible to assign a set of thresholds to multiple components at one shot?**

  Yes. By defining threshold rules and creating threshold groups, you can actually assign a set of pre-defined thresholds to multiple components, simultaneously. Please refer to the *eG User Manual* for further information.

- **What are Global Thresholds?**

  While monitoring large environments, some tests executed by the eG agent report statistics on hundreds of descriptors. For example, the UserProfile test reports the profile size of each and every Citrix or Terminal server user of a server. Likewise, the WinSvcStatus test reports on the availability of each and every service of a Windows system. For such tests, storage of the threshold values for each hour for each descriptor can result in significant disk space usage in the eG database. In order to enable administrators to optimize database usage for tests that do not use the automatic threshold computation (i.e., relative thresholding) capability, eG Enterprise offers the **FIXED THRESHOLD CONFIGURATION** page. To access this page, select the **Global Thresholds** option from the **Threshold** sub-menu of the **Agents** menu.

- **Is there a separate interface for viewing threshold configurations?**

  Yes. Follow the menu sequence, Agents -> Thresholds -> View, to view the threshold configurations of chosen components/component-types.

6.16 Configuring Alarm Policies

- **What are the alarm policies currently supported by eG? How are they different from each other?**

  The eG Enterprise Suite includes four preset alarm policies namely **standard**, **immediate**, **shortterm**, and **longterm**. These alarm policies are used by the system to determine when to generate alarms. For instance there might be an instantaneous spike in the CPU utilization of a system. While an instantaneous spike may not be a problem, a set of periodic spikes or a persistent increase in the measurement may indicate a problem. To differentiate between these different scenarios, eG uses two parameters **Window size** and **Number of crossings**. While Window size reflects the number of measurements in the past that are used to determine the current state of a measure, the number of crossings indicates the minimum number of measurements in the window which if abnormal change the state to bad. For eg., if Window size is 3 and Number of crossings is 2, then it means that eG will generate an alarm if 2 readings out of every 3 readings of a measure are abnormal.

- **How do I configure a new alarm policy?**

  Administrators can configure new alarm policies by choosing the **Alarm Policy** option from the **Agents** configuration menu. The existing set of alarm policies will be displayed. New policies can be added using the **Add New Policy** option.

- **Can I modify the settings of an existing alarm policy?**

  Yes. To modify an existing alarm policy, choose the **Alarm Policy** option from the **Agents** configuration menu. The existing set of alarm policies will be displayed. Click the **Modify** button beside the alarm policy to be modified.

- **Can I delete an existing alarm policy?**

  Yes. Choose the **Alarm Policy** option from the **Agents** configuration menu to view the existing set of alarm policies. Click the **Delete** button beside the alarm policy to be deleted.

6.17 Viewing the Agent Status

- **How can I know whether an agent is running?**

  The Agents -> Status menu provides the status of a selected agent-type. Agents are of 3 types, namely: Basic, Premium, and External.

  a. **Basic Agent:** A **basic** agent can be used to monitor only the operating system of a host and the processes running on it. To use a basic agent, the user must manage the host as a *Generic*, *SNMP Generic*, *AS400*, *Netware*, or an *Event Log* server.

  b. **Premium Agent:** If any applications on the host (e.g., Web, email, DNS, etc.) have to be monitored, the agent is counted as a **premium** agent. An external agent is regarded as a **premium** agent.

  c. **External Agent:** These are agents that execute tests from locations external to the servers and network components of an infrastructure.

  Selecting any of the above agent types will invoke a list of agents of that type accompanied by their status.
  A symbol against each agent indicates that the agent has been deployed. A against each agent implies that the agent has not been deployed. A full charge on the cell in the **Status** column indicates that the eG agent is up and running.

- **Can I restart all agents from the eG administrative interface?**

  Yes. Clicking on the **Restart All Agents** button in the **AGENTS STATUS** page (Agents -> Status -> Details menu sequence), remotely initiates an agent restart.

- **Can I restart one particular agent?**

  Yes. To do so, first, click on the **RESTART** button corresponding to the agent in the **AGENTS - STATUS** page.

- **What does the INFORMATION page reveal?**

  This page reveals agent-related information, which includes:

  - The **Agent IP/Nickname**
  - An indicator as to whether the auto upgrade capability has been **Enabled** for that agent, or **Disabled**
  - The ID of the last upgraded package (if any) (if no upgrading has occurred, then this will be 'None")
  - The date and time at which the agent was last upgraded

- The **HostName** of the agent
- The operating system on which the agent is executing
- The current version of the agent
- The date and time at which the agent last updated the manager with configuration changes
- A **RESET** button

- **Can I enable output logging for a particular agent from the eG administrative interface? How do I view the logs?**

  Yes, you can. For this, set the **OUTPUT** flag that corresponds to an agent in the **AGENTS-STATUS** page to **ON**. To view the log, click on the **LOG** icon that corresponds to that agent.

- **What is the significance of the RESET button in the INFORMATION page?**

  Once an agent is upgraded, information regarding the upgraded package will be registered with the manager. The **INFORMATION** page provides that information only. Now, the next time the agent requests for an upgrade, the manager checks whether any newer upgrades are available. If any such upgrade is found, it sends the same to the agent. If for some reason the information pertaining to the last upgrade has to be cleared from the agent's upgrade history, then click on the **RESET** button. This ensures that the details of the last upgrade are lost, and helps the agent download the last upgrade once again from the manager.

- **What is an Upgrade? Do I need to manually upgrade every eG agent?**

  No. Upgrades (or patches) to the eG agents add new features and enhancements to the eG product suite. Manual installation of the agent upgrades involves a lot of time, labor and cost, especially in environments comprising of hundreds of agents spanning multiple locations. In order to simplify the process of deploying the agents, eG Enterprise offers the auto upgrade capability. By default, this capability is disabled for all agents. Once it is enabled, then, the next time the agents check the manager for the existence of an upgrade, the manager will send the upgrade (if any) to the agents. The agent will then install the upgrade automatically.

- **How do I enable the auto upgrade capability of the eG agent?**

  Please refer to the *eG User Manual* for a detailed procedure on how to enable the auto upgrade capability.

6.18 Manager Redundancy

- **What is a redundant eG manager configuration?**

  In the default deployment, the eG Enterprise Suite has a single central manager that receives performance metrics from the agents, detects anomalies, and sends email and SMS alerts to administrators regarding a probable problem condition. In this scenario, if the eG manager fails, critical issues in the target infrastructure (e.g., failure of a network router, stopping of a key application process) will go unnoticed and therefore, unattended. To ensure high availability of the eG monitoring solution, the eG Enterprise Suite offers a redundant manager option wherein a secondary manager can act as an active or passive standby for the primary manager. This

capability, together with the ability to deploy redundant external agents in multiple locations, ensures that there is no single point of failure for the monitoring solution.

- **What is a Manager Cluster? Differentiate between the *Active-Active* and *Active-Passive* clusters.**

  A cluster can have a single **primary manager** and one **secondary manager**. An *Active-Active* cluster is one where both the primary and secondary managers can both have agents reporting measures to them during normal operation. Alternatively, administrators can opt for an *Active-Passive* setup wherein the primary manager alone will manage all the agents deployed during normal operation. The secondary manager in such a setup will remain passive - i.e., will be up and running, but will not manage any agents unless the primary manager fails.

- **Is the redundant manager capability license controlled?**

  Yes. The Manager Redundancy is a license-controlled feature of the eG Enterprise Suite, and is governed by the **Cluster Type** condition in the eG license. If the **Cluster Type** condition contains the value *Not Supported*, it indicates that the current installation of the eG Enterprise Suite supports a single manager only. If **Cluster Type** is set to *Active-Active* or *Active-Passive*, then it indicates that manager redundancy has been enabled for that eG installation.

- **By default, all agents report to the Primary Manager in a cluster. Can this default setting be overridden? If so, how?**

  Typically, the manager to which an agent reports will be set during the time of agent installation itself. To facilitate manager redundancy, this manager-agent association can be overridden using the eG administrative interface. To change the mapping, use the **ASSIGN AGENTS** page (Agents -> Redundant Manager -> Assign Agents menu sequence) of the eG administrative interface.

- **Is the option to manually override the manager-agent association available to all types of clusters?**

  No. The option to assign agents to specific managers will not be available for an *Active-Passive* cluster.

- **What happens when a manager that was down earlier comes back up?**

  When a manager comes back online, the agents which had earlier switched to report to the other manager in the cluster will automatically begin reporting measures to the original manager to which they were assigned. Also, when a manager is offline, the other manager receiving data from agents will store the data that it receives from the agents locally. When the manager comes back up, the other manager will transmit the saved data to the manager that has just come up. This ensures that there is minimal data inconsistency between the different eG managers.

- **What do I do to remove a primary manager from a cluster?**

  To remove a primary manager from a cluster, the secondary manager should be converted into a primary manager. For further details about configuring manager redundancy, refer to the *eG User Manual*.

- **Do all managers in a cluster allow admin logins?**

  No. An administrator can login and make configuration changes using the primary manager only. To simply monitor the environment without making any configuration changes, the administrator can login to the secondary manager as well. Monitor users can login to any of the managers in a redundant *Active-Active* cluster. In an *Active-Passive* cluster though, monitor users can login to

the primary manager only as long as it is available. If primary is down, they will be able to login to the secondary manager.

6.19 Audit Logging

- **What is an Audit Log?**

  An audit log can be best described as a simple log of changes, typically used for tracking temporal information.

- **Why should the eG manager create and maintain audit logs?**

  The eG manager can now be configured to create and maintain audit logs in the eG database, so that all key configuration changes to the eG Enterprise system, which have been effected via the eG user interface, are tracked.

  The eG audit logs reveal critical change details such as what has changed, who did the change, and when the change occurred, so that administrators are able to quickly and accurately identify unauthorized accesses/modifications to the eG Enterprise system.

- **How to enable audit logging?**

  By default, audit logging is disabled. To enable the capability, follow the steps given below:

  - Login to the eG administrative interface.
  - Open the **SETTINGS** page using the menu sequence: Configure -> Settings.
  - Expand the **Manager Settings** tree in the page and click the **Auditing** node tree.
  - In the **AUDITLOG** section that appears in the right panel, set the **Enable auditlog** flag to **Yes**.
  - Click the **Update** button to save the changes.

- **How to generate AUDIT LOG REPORTS?**

  To view the details logged and analyze their implications, eG Enterprise provides an exclusive **Audits** menu in its administrative interface, using which you can generate a variety of **AUDIT LOG REPORTS**.

Chapter 7

Infrastructure Monitoring using the eG Enterprise Suite

This chapter provides solutions to the queries that may arise while monitoring IT infrastructures using the eG Enterprise Suite of products.

- **Does the eG Enterprise Suite need time synchronization between the manager and agents?**

  No. eG does not require explicit time synchronization between systems in the target environment. The eG agents and manager periodically synchronize their notions of time. The agents report all the measurements based on the manager's time settings.

- **How can I access the eG reports?**

  eG supports a web-based user interface. Hence, users can access the eG reports from a browser.

- **What does the UNKNOWN state mean?**

  The unknown state depicts a condition when the eG Enterprise Suite is not able to make a valid measurement. This may be because the agent is not running. Alternatively, if the component to be measured is not running (e.g., a WebLogic application server is not running), then also the eG Enterprise Suite may not be able to make valid measurements (because the eG Enterprise Suite may depend on this application to provide the measurements).

- **How often does thresholding and trending take place?**

  Thresholding runs once an hour. Trending is done around midnight each day.

- **How will I know if the eG manager detects a problem and how soon?**

  The time required for the eG manager to detect a problem depends on threshold policy and the frequency of tests. As soon as it detects a problem in the environment, the eG manager sends alerts through email within the next 5 minutes.

- **How will I know if the eG agent is not running?**

  No measures would be reported if the eG agent is not running. This will result in the measures appearing in Blue (unknown state) on the user interface.

- **How can I view the entire list of alarms that were generated?**

  Select **History** from the options available under **Alarms** in the menu running across the top of monitor pages to view the entire list of alarms pertaining to the entire infrastructure over a period of time.

- **What do I do to view only the current set of alarms?**

  Choose **Current** from the options available under **Alarms** in the menu running across the top of monitor pages to view the current set of alarms.

- **Can I acknowledge an alarm?**

  If your profile authorizes you to acknowledge alarms, then **Yes**, you can acknowledge an alarm.

- **How do I acknowledge an alarm?**

  Select the check box corresponding to the alarm in the **CURRENT ALARMS** window and click the **Acknowledge** button. Provide a brief description of the acknowledgement and submit it by clicking the **Submit** button therein.

- **How does alarm acknowledgement help?**

  By acknowledging an alarm, a user can indicate to other users that the issue raised by an alarm is being attended to. In fact, if need be, the user can even propose a course of action using this interface. In such a case, a user with **Admin** or **Supermonitor** privileges (roles) can edit the acknowledgement by providing their own comments/suggestions on the proposed action. The acknowledgement thus works in three ways:

  - Ensures that multiple members of the administrative staff do not unnecessarily invest their time and effort in working on a single issue;
  - Serves as a healthy forum for discussing and identifying permanent cures for persistent performance ills;
  - Indicates to other users the status of an alarm

- **Can I delete alarms?**

  If your profile authorizes you to delete alarms, then **Yes**, you can.

- **How do I delete alarms?**

  Select the check box corresponding to the alarm to be deleted in the **CURRENT ALARMS** window and click the **Delete** button. Provide a reason for the deletion and click the **Submit** button.

- **Can I view the history of alarm acknowledgements and deletions?**

Yes. The eG monitoring console provides you with many interfaces with the help of which the complete history of acknowledgements and deletions can be viewed.

Alarm Acknowledgement History

In large environments, it is but natural that the same set of components are assigned to multiple users for monitoring. In such environments, some/all the users with monitoring rights to a component might want to post their comments for an alarm related to that component. If acknowledgment rights are granted to all these users, then each of them can login to the monitor interface and provide an acknowledgement description for the same alarm.

eG Enterprise maintains a history of the acknowledgement descriptions provided by multiple users with rights to monitor a single component, and lists the entire history the next time one of these users attempts to view the acknowledgement details in the **CURRENT ALARMS** window. To view the acknowledgement history, move your mouse pointer over the **A** icon that appears alongside any acknowledged alarm in the **CURRENT ALARMS** window. This way, the administrative staff can share the responsibility for resolving critical issues, brainstorm online to identify accurate remedies, and even provide each other with quick updates on problem status.

If a past alarm is associated with one/more acknowledgement descriptions, then the **EVENT HISTORY** page too will automatically display the acknowledgement history below that alarm.

Sometimes however, to perform better problem diagnosis, you might want to review specific acknowledgment descriptions associated with an alarm and not all of them. For instance, while two users - *john* and *elvis* - may have acknowledged an alarm raised on an Oracle database server, you might only want to view user *john's* acknowledgement description. To facilitate such selective viewing of acknowledgement information, eG Enterprise provides a dedicated **ACKNOWLEDGEMENT HISTORY** page; this page provides a wide variety of filter options with the help of which you can quickly and easily run a search across all alarm acknowledgements, and swiftly locate the acknowledgment information of interest to you. To access the **ACKNOWLEDGEMENT HISTORY** page, follow the *Alarms -> History of Acknowledgements* menu sequence in the eG monitoring console.

Alarm Deletion History

The **EVENT HISTORY** page has also been embedded with the intelligence to indicate whether a past alarm was deleted or not. If a past alarm had been acknowledged, then this page will automatically display the acknowledgement history below that alarm. If that alarm had been deleted at some point of time in the past, then the history will display an additional entry providing the details of the deletion. These details include: the user who deleted the alarm, the reason for the deletion, and the date and time of deletion.

Sometimes, you might want to view the details of all alarms that were deleted by a particular user. Similarly, you might want to view only the details of those alarms that were deleted during the last 24 hours. Since the **EVENT HISTORY** page allows you to search based on general alarm information alone and not on deleted alarms, this page cannot be used for performing the search operations mentioned above. To zoom into the details of specific deleted alarms therefore, eG Enterprise offers a dedicated **DELETION HISTORY** page. This page provides a variety of filter options using which you can quickly access the alarm deletion details that you may require. To access this page, follow the menu sequence: *Alarms -> History of Deletion*.

- **What does the Monitor Home page display?**

This page quickly updates the monitor user with the health of the entire monitored environment. The page reveals the following information:

- The first section is the **Current Status** section that reveals at a glance, the status of the measurements reported to the eG manager. Besides displaying the total number of monitored components and the number of performance metrics collected by the eG agents from these components, this section also reveals the percentage of total measurements that are in the critical, major, minor, normal, and unknown states. Using this information, an accurate assessment of the overall infrastructure performance can be made. You can click on any measurement state displayed in this section to view the list of open alarms that correspond to that state.
- Below the **Current Status** section, is the section that reveals the **Infrastructure Health**. Since the health of an infrastructure depends entirely upon the performance of each of its key ingredients - namely, the Components, Zones, Services, and Segments - this section takes the help of a bar graph to clearly indicate the number of zones, services, segments, and components that are in the Critical, Major, Minor, Normal, and/or Unknown states.

  The table below explains the color-coding scheme adopted by eG for indicating the states of the services/segments/components:

Color	State
Red	Critical
Orange	Major
Pink	Minor
Green	Normal
Blue	Unknown

If you click on a division in the **Zones** bar graph, you will be lead to a page that lists the zones which are in that particular state.

Clicking on a division in the **Services** bar graph in the **Infrastructure Health** section will lead you to a page that lists the services, which are in that particular state. By default, against every service in this page, the service components (i.e., the components that are associated with the service) and their current state will be displayed

Clicking on a division in the **Service Groups** bar graph will lead you to a page that lists the service groups, which are in that particular state.

Next, to know how well every segment configured in the environment has fared, take a look at the **Segments** bar graph in the **Infrastructure Health** section. Clicking on a division in this graph will enable you to view the list of segments in a particular state. Alongside each segment, the IP/hostname of the segment components (i.e., components that are part of the segment) and their current state will be displayed.

Finally, click on a division in the **Components** bar graph for a quick update on components that are currently in a particular state.

- Below the **Infrastructure Health** section you will find a **Measures At-A-Glance** section that provides the min/max values of critical measurements updated in real-time. By default, using this section, you can quickly find answers to the following critical performance queries:

  (a) Which host across the environment is consuming the maximum CPU?

  (b) Which host in the monitored infrastructure has very little free memory to its credit?

  (c) Which disk partition on which host is utilized the maximum?

  (d) Which is the Citrix server that supports the maximum number of active sessions?

  (e) Across all monitored Citrix servers, which application is CPU-intensive and which Citrix server is it executing on?

  (f) Which is the web server that services the maximum number of requests over time?

  (g) Which network interface is using up a lot of bandwidth?

  (h) On which host are TCP retransmits very high?

  (i) Where in the target environment is network latency the maximum?

  (j) TCP connections to which port are taking too long?

  (k) Which host is currently not available over the network?

  You can however, change this default setting to reveal more or less - i.e., you can add to the measure list displayed here, or remove a few measures from the displayed list. To achieve this, do the following:

  - Login to the eG administrative interface as *admin*.
  - Go to the **MONITOR SETTINGS** page using the menu sequence, Configure -> Monitor Settings.
  - Click on the **Measures At-A-Glance** button in the **MONITOR SETTINGS** page.
  - The resulting **MEASURES AT-A-GLANCE CONFIGURATION** page will display the default measure configurations for the **Measures At-A-Glance** section. To know how to manipulate the controls in the page to add more measures or remove a few of the pre-configured measures, refer to the *eG User Manual*.

  The min/max values of the configured measures will then be displayed in the **Measures At-A-Glance** section of the Monitor Home page. The first column of the **Measures At-A-Glance** section indicates the current state of the measure (whether Normal/Critical/Major/Minor/Unknown). The **Measure** column is where the configured measures will be displayed. Similarly, each of the configured tests will appear in the **Test** column. Besides, a **Server** column exists, which displays the name of the component and the descriptor, which has currently registered the maximum/minimum value (as the case may be) for every chosen measure. Finally, the current value of the measures for the displayed components will be displayed in the **Value** column.

  Optionally, you can even switch off the **Measures At-A-Glance** section. To do so, you will have to set the **Compute top metrics** parameter in the **MEASURES AT-A-GLANCE CONFIGURATION** page of the eG administrative interface to **No**. By default, this parameter is set to **yes** indicating that, by default, the monitor home page will contain the maximum/minimum computations for

measures (i.e., the **Measures At-A-Glance** section). When it is set to **no**, the **Measures At-A-Glance** tab is hidden from the home page, and the **Event Analysis** tab alone appears.

- The **Event Analysis** tab page, when clicked, lists the top-5 layers that were most affected by performance issues.

  Corresponding to every layer name in the **Event Analysis** section, you will see the number of alarms that are currently open for that layer, the average duration of the open alarms, and the maximum duration for which an alarm had remained open. To the right of the tab page, you will find a list of components that experienced the most number of performance issues during the last one hour (by default). Against every component listing, the corresponding component-type, and the number of events the component encountered during the default period of one hour, is displayed. This information brings to light the most problem-prone components in the environment. Clicking on a component/component-type in this list, reveals the layer model, tests, and the last set of measurements that the eG agent reported for that component. To analyze events across a broader time window in the past, select a different timeline from the **Components with most events in the last** list box here.

  The details available in the **Event Analysis** section serve as an effective indicator of the efficiency of the administrative staff in resolving performance issues. To view the complete history of alarms in the environment, click on the **Click here for more events >>** link.

- The **Components At-A-Glance** section comprises of a bar graph depicting the number of components of each type that are being monitored, and their respective states. Clicking on a bar will take you to a page that lists the individual components of the corresponding type and their current state.

- **What happens when I click on a zone in the ZONES LIST page of the eG monitoring console?**

  The **Zone Dashboard** appears that provides a quick overview of the performance of that zone. For more details on the dashboard, refer to the *eG User Manual*.

- **Can I view the location of zones configured?**

  Typically, zones are associated with different geographies. While monitoring large infrastructures therefore, eG Enterprise allows you to drill down to view the exact geographic area where a zone operates, and instantly evaluate the performance of the different zones spread across the different locations worldwide. To access the map interface that provides this visual treat, select the **Map** option from the **Zones** menu in the eG monitoring console. The page that then appears indicates the zone locations and their current state. The geographic display of the maps within the eG Enterprise console is achieved through the integration of the eG management console with Google maps.

- **What is the zone tree-view?**

  To the left of the zone map, you will find a **TREE VIEW** section displaying a zone tree. While the map enables you to determine the location and the current state of the zones, using the tree view, you can quickly determine the following:

  - The names of the zones for which locations have been explicitly defined in the eG administrative interface
  - The state of each of these zones

- The names of sub-zones added to a zone, provided the location of the sub-zone has been configured

➢ **What do I do to view the status of all managed applications in a target environment?**

Just select the **Servers** option from the **Components** menu in the eG monitoring console.

➢ **What do I do to view the status of all managed systems in a target environment?**

Just select the **Systems** option from the **Components** menu in the eG monitoring console.

➢ **Why is a 'System View' necessary?**

While eG Enterprise focuses primarily on monitoring applications, many administrators still prefer to view their infrastructure from a hardware perspective – i.e., as systems they support. The eG monitoring console therefore provides a "system view", which represents the overall health of systems in the target infrastructure, with a mapping of the applications that are executing on these systems.

➢ **How do I determine the health of virtual hosts?**

If your environment comprises of virtualized components such as VMware® ESX hosts or Solaris virtual servers, then selecting the **Servers** option from the **Virtual** sub-menu of the **Components** menu will allow you to view the current state of the that have been managed by the eG Enterprise system

➢ **How do I access the VM Dashboard?**

The VM dashboard can be accessed by following the menu sequence, Components -> Virtual -> Dashboard, in the eG monitoring console.

➢ **What are the benefits of the VM Dashboard?**

Using this dashboard, administrators can:

- Detect from a mere glance, excessive resource usage by any VM, resource pool, or a physical server from across the environment, regardless of the virtualization technology in use; quickly diagnose the root-cause of the resource drain with the help of efficient drill down features;
- Accurately identify resource-intensive VMs/resource pools/physical servers;
- Instantly spot a powered-off VM anywhere in the environment;
- Know which VM is currently operating on which physical server, and thus keep tabs on VMotion/XenMotion (as the case may be) activities;
- View a consolidated list of issues currently encountered by physical servers and virtual machines across the environment and also per virtual component, so as to ease the troubleshooting efforts of a dedicated help-desk;

➢ **What are the contents of a VM dashboard?**

The **VM Dashboard** of comprises of two panels. The left panel boasts of a tree-structure, comprising of a global node named **Zones**. All the zones/farms in the target environment that have been configured with one/more virtual hosts (eg., VMware ESX servers, Citrix XenServers, etc.), will be

the sub-nodes of the **Zones** node. If you expand a particular zone node in the tree, you will find that the virtual component-types that have been added to that zone appear as its sub-nodes. Independent virtual hosts that are not part of any existing zone will be automatically grouped under the **Default** zone tree. If you expand a virtual component-type under a zone, then all virtual hosts of that type will appear as its sub-nodes. Expanding a virtual host tree will reveal the VMs and resource pools that are executing on that tree. Similarly, if you drill down a resource pool in the tree, the VMs available in the pool will be listed as its sub-nodes. The state of a node in the tree depends upon the current state of its sub-nodes.

The right panel is a context-sensitive panel, the contents of which will vary according to the node chosen from the left panel. By default, this panel provides 3 tab pages – a **VMs** tab page that provides current status updates related to virtual machines, a **Hosts** tab page that displays virtual host-specific metrics gathered in real-time, and a **Current Events** tab page that lists the problems currently experienced by virtual hosts and guests. By default, all the tab pages provide information pertaining to managed *VMware vSphere ESX servers* in the environment. Accordingly, the **VMware vSphere ESX** option is chosen from the drop-down list in the top, right corner of the right panel. You can view details for a different virtual component type, by selecting a different option from this drop-down list.

- **Are dashboards available only for Zones and Virtual Infrastructures? What about other servers/applications?**

The eG monitoring console now provides dashboards for each of the critical layers of an application/device that, you can zoom into the current and historical performance of that entity using a single interface. Dashboards are also available for segments and services.

- **When you have the layer model, why do we need application-specific dashboards?**

By default, the primary means of analyzing the performance of an infrastructure component – a network device, a server, or an application – is by using a layer model representation. The layer model is hierarchically structured and each layer mapped to specific functionality of the component. Tests and measurements are mapped to each layer and the state of a layer is determined based on the status of the measurements mapped to it. The layer model representation is used for automatic correlation of metrics – when a similar priority problem happens at two layers, the problem at the higher layer of the layer model is attributed to being caused by the lower layers. The layer model representation has several advantages:

- By using a common model for representing heterogeneous infrastructure components, eG Enterprise makes it easy for an administrator to monitor different components with diverse functionality (the monitoring model in the eG Enterprise representation was similar).

- The layer model representation makes it easy to demarcate problems – e.g., is a problem with an application being caused in the operating system layer, or in the network layer, or in the application layer?

The main limitation of the layer model representation is that if an administrator is interested in an at-a-glance view of the key metrics for a component, this is not available. Further, when looking for a specific measurement, administrators need to know which layer the measurement maps to. Otherwise, they would have to click through each of the layers to find the measurement of interest.

To address these shortcomings, eG Enterprise now includes specialized dashboards for network, system, and application monitoring. In addition to the layer model, administrators can access key metrics at the network, system, and application layers directly from the dashboards. For example, in the system dashboard, administrators can view at a glance, the CPU utilization, memory utilization, disk usage, current system configuration, top CPU and memory consuming processes, and other key system metrics.

Thus, these dashboards:

- Serve as a single, central console that not only depict the current state of a layer, but also instantly indicate the root-cause of issues pertaining to that layer, thereby enabling administrators to go from problem effect to the problem source in no time!
- Combine both raw and graphically represented data, and facilitate an in-depth analysis of not just live performance, but also the historical performance of a particular layer, thus shedding light on potential anomalies;
- Aid administrators in effectively analyzing the past trends in the performance of a layer, so that they can easily forecast future performance;
- Enable service level audits on-the-fly, and thus help administrators accurately determine when a layer slipped from the desired performance levels.

- **Briefly describe the service dashboard?**

Some mission-critical environments may support services, which are delivered by a large number of components that share intricate inter-relationships with one another. Administrators of such environments may find it tedious to look across multiple interfaces - eg., the **SERVICE LIST** page, the **SERVICE TOPOLOGY** page, the **Layer Model** tab page, etc. - to identify problematic services and to isolate the problem source. Such administrators would therefore prefer a single, central, **Service Dashboard** that not only introduces them to the problematic services, but also readily provides them with all the information that they may require (such as the components associated with the service, the relationship between the components, the state of each component, and the metrics they report) to deep dive into the issue at hand and zero-in on its root-cause, without requiring cumbersome navigation. eG Enterprise provides such a dashboard.

The **Service Dashboard** boasts of a left panel that contains a tree-structure, with a global **Services** node. Expanding this node will list all the services that are monitored in the target environment as the sub-nodes. When you expand a particular service node in the **Services** tree, the components that are associated with that particular service will appear as its sub-nodes. The current state of every service component will also be indicated.

When you click on a particular service sub-node, a context sensitive right panel will appear with contents varying according to the service chosen from the left panel. By default, the right panel comprises of 3 tab pages – a **Systems** tab page providing an at-a-glance view of all host level measures for each component engaged in the delivery of the chosen service, a **Components** tab page providing an at-a-glance view of all the application level metrics for each service component, and the **Topology** tab page displaying the service topology. An additional **Transactions** tab page will appear for **web site services alone**. The **Transactions** tab page will display all the transactions that have been configured for a chosen web site, the current state of each transaction, and the metrics collected in real-time for every transaction.

- **What dashboards accompany each application/device that is monitored, and what is their purpose?**

  While all components monitored come with a **System** and **Network** dashboard, **Application** dashboards are available for a few key applications only.

  The **System Dashboard** of an application allows you to focus on the performance of the operating system on which that particular application runs - i.e., the **Operating System** layer of an application. Using the **System Dashboard**, administrators can determine the following:

  - The current status of the application host;
  - The problems that the host is currently facing, and the type and number of problems it encountered during the last 24 hours;
  - The current system configuration (if the eG license enables the **Configuration Management** capability);
  - The current state of the critical parameters related to system performance;
  - How some of the sensitive performance parameters have performed during the last 1 hour (by default);
  - The resource-hungry processors supported by the host, and the disk partitions on the host that are currently experiencing a space crunch.

  The **Network Dashboard** allows you to zoom into the performance and problems pertaining to the **Network** layer and related layers of a target application/device. Using this dashboard, you can:

  - Determine whether/not the application/device currently experiences / has in the past experienced network-related issues;
  - Accurately identify the network parameters that are currently failing;
  - Understand the current network configuration;
  - Analyze network performance over time, study the trends in network connectivity and usage, and accurately deduce problem/performance patterns.
  - Identify persistent problems with network health and the network-related layers responsible for the same;

  In order to ascertain how well an application is/has been performing, analysis of the performance of the **System** and **Network** layers of that application alone might not suffice. A closer look at the health of the **Applicaton Layers** is also necessary, so as to promptly detect instantaneous operational issues with the target application, and also proactively identify persistent problems or a consistent performance degradation experienced by the application. To provide administrators with such in-depth insights into overall application performance and to enable them to accurately isolate the root-cause of any application-level slowdown, eG Enterprise offers the **Application Dashboard**.

- **Can the contents of a dashboard be modified?**

  Yes. **Admin** and **Supermonitor** users are alone authorized to modify the contents of a dashboard by clicking on the button at the top of the dashboard. A **Dashboard Settings** window will then appear using which such users can alter the default dashboard settings.

- **What are the applications for which Application Dashboards are provided?**

  Currently, only the following applications support **Application Dashboards**:

  - Java application
  - MS SQL server
  - Citrix XenApp server

- **What about other applications for which an Application Dashboard is not offered out-of-the-box?**

  Though the contents of the system/network/application dashboards are customizable, the right to customization rests with the **Admin** and **Supermonitor** users alone, and not all users to the eG monitoring console. This means that if any monitor user logs into the console, he/she will only be allowed to use the pre-defined dashboards. These dashboards, as we know, will focus on only those metrics that an **Admin** or **Supermonitor** has configured - a normal monitor user can neither customize the layout nor alter the contents of such dashboards.

  Also, note that by default, dashboards are available only for those applications that are supported out-of-the-box by the eG Enterprise Suite. For in-house/legacy applications that may have been integrated into the eG Enterprise Suite using the Integration Console plugin, ready-to-use dashboards are not available.

  Therefore, to enable every user with monitoring rights to personalize his/her dashboard experience, the eG monitoring console allows the creation of **Custom Dashboards**. These dashboards can be designed for both existing applications and legacy applications. This capability empowers users to control what data is to be displayed in the dashboards and how to present it (whether to use dial charts or digital displays or comparison graphs or tables). This way, users can see what they want to see in the dashboards.

  You can build custom dashboards for a particular application or build a generic dashboard that displays the status of and metrics collected from a multitude of applications.

- **What are the different types of graphs that I can view via eG user interface?**

  eG Enterprise includes a variety of graphing capabilities for manual diagnosis. eG currently supports three kinds of graphs. A measure graph is used by a user to plot the instantaneous value of any of the measurements made by eG with time of day. The measure graph can be accessed via the **MEASURE** option under **GRAPHS** in the menu across the top of the monitor pages. A summary graph gives an overall picture of the percentage of good, bad and unknown measurements over a period of time. Since Internet traffic is very bursty, using the measurement graphs over a long time window (greater than a week) to view trends in the measurements can be very difficult. To make it easier to analyze trends, eG monitors and stores trend data on an hourly, daily, and monthly basis. The eG user interface allows the trend data to be plotted on a web browser.

- **What are the different types of trend graphs that can be generated using the eG monitoring console?**

  By default, a trend graph takes the minimum and the maximum value of measurements over a period of time into consideration. Accordingly, the **Graph** list displays **Min/Max** by default. Alternatively, you can even generate a trend graph where the **average** values of a chosen measure are plotted over a period of time. To achieve this, simply select the **Average** option from the **Graph** list in the **TREND GRAPH** page. For instance, you can plot a trend graph that depicts how many TCP

connections on an average were established with a critical Terminal server, every day during a couple of weeks; besides indicating the normal load on the Terminal server, such a graph also enables you to understand whether the Terminal server has been adequately tuned to handle higher loads, and thereby helps you make effective sizing recommendations for the future. Likewise, you can also choose **Sum** as the **Graph** type to view a trend graph that plots the sum of the values of a chosen measure for a specified timeline. For example, a **Sum** graph of TCP connections to a Terminal server serves as an accurate indicator of how busy the Terminal server was during a given period.

Note:

The capability of the eG manager to compute the **Sum** and **Average** of metrics is governed by the **Compute average/sum of metrics while trending** flag in the **MANAGER SETTINGS** page (Configure -> Manager Settings) in the eG administrative interface. By default, this flag is set to **Yes** indicating that eG Enterprise computes and stores the average and sum values of every performance metric in the database, by default. If, for some reason, you want to disable this capability, just set this flag to **No**, and **Update** the changes.

- **I need to know which Citrix server-related metrics are currently under-performing, but I really don't wish to take the cumbersome layer-test-measure route. Do you have an alternative?**

eG agents are capable of extracting a wealth of performance information pertaining to the managed components. Larger the environment, larger will be the number of measures collected. The biggest challenge for the administrators of such large environments therefore, is to isolate and attend to the anomalies surrounding certain critical performance metrics. For instance, administrators of large, mission-critical Citrix environments might want to focus on some sensitive performance areas such as user sessions to the Citrix farm, resource usage by applications published on the farm, the rate of growth of the user profiles, etc. The metrics related to these areas are mapped to different layers and different tests of the Citrix monitoring model. Instead of taking the longer layer-test-measure route, administrators might prefer a single interface that provides a consolidated list of metrics collected from across the Citrix environment, so that they can easily pick and choose the metrics of interest to them. The **MEASURES** page that appears upon clicking the **Measures** menu option provides this much-needed measure focus.

- **Can I maintain an information base for the different problems encountered in the environment?**

The users can keep track of the problems arising in their environment and the solutions to these problems by making an entry in the feed back form that can be obtained on clicking the **FEEDBACK** option available beside the last measurements made. This page enables a user to record the details of how specific problems were identified and how the problems were fixed. This information is maintained in the eG database and represents a location-specific knowledge base that can be queried by monitor users of the eG Enterprise Suite.

- **I have a problem in my environment? Can I know if similar problems were fixed before?**

As mentioned above, the eG database also keeps a repository of the problems encountered in the past and how they were fixed based on the input given by the users via the feedback form. This knowledge base that can be queried by monitor users of the eG Enterprise Suite when a problem and determine how a similar problem had been fixed in the past. The history of problems can be obtained by clicking on the **HISTORY** button available beside the last measurements made.

- **How do I interpret a measurement?**

  For details on how to interpret a measure, kindly utilize the online help feature of the eG user interface. *The eG Measurements Manual* also provides details regarding the interpretation of the measurements.

- **The *CPU utilization* measure of my SystemDetails test is in Critical state with a value of nearly 98%. I would like to know which process(es) is consuming this much Cpu time. Can the eG Enterprise Suite help me in this regard?**

  Yes. To know the process(es) consuming more CPU time, you will just have to click the **DIAGNOSIS** button (resembling a magnifying glass) available against the *CPU utilization* measure in the **COMPONENT MEASURES** page. This will result in a page that displays the top 10 processes consuming more CPU time.

- **Is there an alternative means of viewing the detailed diagnosis of a measure?**

  Yes. The **DETAILED MEASURES** page can also be accessed using the menu sequence: Options -> Detailed Diagnosis. In the page that appears next, you can select the **Service**, **Component**, **Layer**, **Test**, **Measurement**, and the period for which detailed measures are to be viewed. Detailed measures that match the specified conditions will then be displayed in the page.

- **I want to compare the disk space usage across all the disk partitions on my host. How do I do this?**

  If a test supports more than one descriptor as in the case of the DiskActivityTest, DiskSpaceTest, etc., then the eG monitor interface allows you to instantly compare performance across descriptors, by clicking on the button that will accompany a descriptor-based test in the layer model page.

- **How is site topology different from segment topology?**

  The segment topology indicates how components are interconnected and which components rely on others for their functioning. The interconnections can represent either physical connections (e.g., a web server connected to a network router) or logical dependencies (e.g., a web server using a web application server). Each interconnection is associated with a direction. The direction signifies cause-effect relationships (if any) between the components being connected together. The segment topology deals only with web servers and not with web sites. Web sites and web servers have a many to many relationship among them. In an ISP hosting scenario, many web sites may be hosted on the same web server. Alternatively, very large eBusiness sites may use multiple web servers to support a single web site. eG supports both of these scenarios. The site topology brings out this relationship.

- **Can operators have different views of the web sites?**

  The eG Enterprise Suite product will enable the users to have different views.

- **Why doesn't the eG manager report only for the period when there was a problem? Why does it report all the measures throughout?**

  The eG Enterprise Suite does relative thresholding based on usage trends and not just problem detection. Hence it requires continuous measures even when the components are functioning properly. This information is also required when people want to look at the usage trends and overall performance for the past.

- **What is one click problem diagnosis?**

  When a monitor user logs into the eG user interface, an alarm window showing all the current alarms in the system will appear automatically. Clicking on an alarm will pin point the measure that is causing the problem in the environment. This is one click problem diagnosis.

- **What is alarm suppression?**

  We will illustrate this phenomenon with the example given below:

  Consider two Oracle databases in the same system. One database is used by one site, and the other database is used by another site. When the disk is running low on space, the DiskTest pertaining to the host layer will generate an alarm for each database. The eG Enterprise Suite is intelligent enough to figure out that the problem is common for all the components residing on this server. Thus, the eG Enterprise Suite will throw this alarm once for all components and not once for every component. This is what is called as alarm suppression.

- **How do I access the eG Reporter?**

  To access the eG Reporter, select the **Reports** option from the **Graphs** menu.

- **What are the different types of Executive Reports that can be generated using the eG Reporter?**

  The eG Reporter offers the following Executive Reports: **COMPONENT**, **SEGMENT**, **SERVICE**, and **COMPONENT TYPE**.

- **What are the different types of Operation Reports that can be generated using the eG Reporter?**

  The eG Reporter offers the following Operation Reports: **NETWORK**, **SYSTEM**, **APPLICATION**, **SITE**, **EVENT ANALYSIS**, and **UPTIME**.

- **How can I save reports for future reference?**

  Once the reports are generated, click on the **FAVORITES** button at the right top corner of the page. Upon clicking, the report conditions specified in the reports page will be displayed. Enter a **Name** for the report, and finally click the **SAVE** button. To recall the report conditions, click the **FAVORITES** tab and then, click on the **FAVORITE NAME** listed therein.

- **How do I schedule reports?**

  The eG Reporter also provides a useful report scheduling capability that automates the process of mailing specific reports (to specific individuals) at pre-defined intervals. To schedule the printing/mailing of the reports displayed a page, click on the **SCHEDULES** button in the tool bar at the right top corner of the reports page. In the page that appears next, provide a name for the schedule against the **Schedule Name** text box. To indicate how often the report would have to be mailed to a specific recipient, select a frequency from the **Mail** list box. The options provided therein include: Daily, Weekly, and Monthly. If the report need not be mailed, then select **None** from this list box. If the report is to be mailed to a specific recipient (i.e., if the **Mail** list box does not contain **None**), then mention the **Mail ID** of the recipient. Then, indicate the **Schedule type**. You can indicate when report scheduling is to occur by picking an option from the **Schedule type** list. To generate schedule reports at the end of every day, pick the **Day end** option from this list. To generate schedule reports at a configured time every day, pick the **Any time** option from this list, and then indicate the exact time of generation using the **Schedule at** time controls that then appear. Finally, click on the **SAVE** button to register the schedule.

- **Can I mail reports directly to specific recipients? If so, how?**

  Yes. The eG Reporter component facilitates mailing of the displayed reports directly to specific recipients. To achieve this, first, click on the **MAIL** button in the reports page. Using the page that appears next, specify the mail IDs of the primary recipients, the mail IDs of the recipients to whom the reports need to be copied, etc.. Also, an appropriate **Subject** can be provided for the mail. Finally, click the **SEND MAIL** to send the reports to the specified recipients.

- **Can I schedule of automatic printing of email reports?**

  By default, the eG Reporter does not allow users to schedule the automatic printing of reports. If you want to schedule report printing, then set the **EnableSchedulePrint** parameter in the **[MISC_ARGS]** section of the **eg_services.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory) to **true** (default is **false**). Once this is done, then, a **Print** list will additionally appear in the **SAVE SCHEDULES** page. Indicate the frequency with which the report is to be printed by selecting an option from the **Print** list box. Here again, selecting **None** would ensure that the report is not printed.

- **By default, how many descriptors can be displayed in the Operations Report page? Is this value configurable?**

  By default, a single graph in an Operations report can display a maximum of 10 descriptors. This value however, is configurable. For that, set the **MaxInfos** parameter in the **[INFOS]** section of the **eg_report.ini** (in the **<EG_INSTALL_DIR>\manager\config** directory) to the required value.

- **What happens when I click on a report with many descriptors?**

  If a report comprising of more than 5 descriptors is clicked on, then an enlarged report will appear providing you with options to restrict your view to specific descriptors only. To know how to use these options, refer to the *eG User Manual*.

- **What are Comparison Reports?**

  Comparison reports enable users to compare, correlate, and analyze across a wide variety of performance metrics generated by the monitored environment, so as to facilitate quick and accurate root-cause identification.

- **What are the different types of Comparison reports offered by eG Reporter?**

  eG offers two types of Comparison reports, namely, **COMPONENTS**, **TESTS**, and **TOP N** reports. The **COMPONENTS** comparison report allows administrators to compare statistics pertaining to selected components. For example, a **COMPONENTS** report can be used to compare the responsiveness of an Oracle server, DNS server, a Web server, and a WebLogic server. The **TESTS** report can be used if any of the metrics collected by a test have to be plotted in the same report to facilitate easy analysis and troubleshooting. For example, a **TESTS** report can be used to compare the CPU utilization of all the hosts in a target environment. Using the **TOP N** reports offered by eG Enterprise, administrators can rank components for every metric collected by the eG Enterprise system. The report also reveals the best/worst players in a particular performance area, thereby bringing to light components that are problem-prone, and which hence require continuous attention. For example, you can view the top 10 consumers of CPU across the infrastructure, using **TOP N** reports.

- **While generating a Components Comparison report, will all the measures chosen for comparison always be plotted in a single graph?**

  No. It is not mandatory for all the measures being compared to be plotted in a single graph. By choosing the **Graph Type** as **Multiple** (instead of the default **Single)**, you can ensure that each of the selected measures are plotted in individual graphs.

- **I had a problem between 3-5pm yesterday. Can I do a post-mortem analysis, find out what was the problem, and how the measures were?**

  Yes. The **Snapshot Reports** offered by the eG Reporter enable you to perform a post-mortem analysis on a monitored environment.

- **What are the different types of Consolidated reports provided by eG Reporter? What does each report type do?**

  The following **Consolidated** reports are offered by eG Reporter:

  - **Server Report**: To simplify individual server monitoring, eG Enterprise offers the **Server Report**, where the critical performance results of a particular server are aggregated over time and made available.

  - **Zone Report**: In large zones that consist of tens of thousands of servers, it is very difficult for administrators to focus on the performance of individual servers - a zone-wide overview of performance is hence desired. Therefore, to enable administrators to evaluate the performance of specific component-types that constitute the zone, the suite provides **Zone Reports**. This report, by default, reveals the average performance of the component-type as a whole (e.g., all Oracle databases), and also the average performance of the individual components of the chosen type. Besides indicating the health of the component-type, this report enables administrators to instantly identify the component(s) that serves as the root-cause of the problems (if any) with the type.

  - **Services Report**: In many environments, there may be other administrators who are only interested in the performance of specific business services. These administrators need reports that are specific to the business services that they are responsible for, and the components involved in supporting these services. Services Reports included in the eG Reporter address this need. There are two sets of reports in the eG Service report category.

    - **WebTransactions Report:** The WebTransactions report is specific to web sites. This report provides a summary of the performance of key web transaction metrics reported by eG agents for each web site. A web site may include multiple web servers, with load balancing being used to distribute the requests across the web servers. For a specific time period (say 2 weeks), a user can use the WebTransactions report to view web transaction metrics across each of the web servers used for a web site. Using this report, a user can determine:

      - Is the transaction load being balanced across all the web servers supporting the web site?

      - Are any unusual error patterns seen on specific web servers (e.g., because some content is not available on one of the servers)?

- Is there any unusual slowness observed on specific web servers?

- **ServicesHealth Report:** Irrespective of whether a service is a web site or not, a user may want to view the resource usage (e.g., CPU, memory, disk utilitization) and other key metrics about all the servers that support a specific service. The ServicesHealth report addresses this requirement. From this report, a user only sees all the components that are specific to a chosen service, and can review performance metrics across all these components to identify any bottlenecks in the infrastructure.

- **IIS Transaction Analysis Report:** Whenever users to your mission-critical web sites (on an IIS web server) or web services (overlying an IIS web server) complaint of slowdowns, the first step to troubleshooting these issues is identifying which specific transaction(s) to the web sites is causing the slowdown. For this, you will have to periodically monitor the responsiveness of key transactions to your web sites, so as to quickly and accurately identify slow transactions. Also, by keenly observing the variations in the responsiveness of transactions over time, you can ascertain how often a transaction experienced slowdowns in a given timeline, determine whether that transaction is prone to delays, and investigate the root-cause of the same. To provide you with such in-depth historical insights into a transaction's performance, eG Enterprise offers the **IIS Transaction Analysis** report.

➢ **Why do we need a comparison report?**

The **Components** comparison report allows administrators to compare statistics pertaining to selected components. For example, consider an environment comprising of a Web server supported by an Oracle server, Ldap server, and a Dns server.

Say that the web server in hosts a web site that depends upon the Oracle server for storage and recovery of critical data. While the web server uses the DNS server for host name resolution, the Oracle server uses the LDAP server.

Assume that the user accessing the web site (on the web server) begins to experience a marked deterioration in the responsiveness of the site. An increase in the response time of a web site could be caused by either/all of the following reasons:

- A network delay between the web server and the user
- A problem in the response time of one/all of the other servers (Oracle, DNS, LDAP) that the web server relies on.
- A comparison report provides an easy way of comparing the response times of each application.

➢ **What is the significance of trend reports and detailed test tables?**

A **Trend** report will not include the data for the current day since trend data is only computed at the end of the day. If the **Trend** option is chosen, the time period of the report should be greater than 1 day. The usage of **Detailed** test tables for generating reports, especially those that span weeks, increases the stress on the eG database, thus resulting in undue delays in report generation. In order to ensure that the database is not choked by such voluminous data requests, you can

configure eG Enterprise to automatically "force" the use of the **Trend** option if the **Timeline** of a report exceeds a pre-configured duration. To specify this time boundary, do the following:

1. Edit the **eg_services.ini** file in the **<EG_INSTALL_DIR>\manager\config directory**.
2. In the **[MISC]** section of the file, you will find a **DetailedTime** parameter.
3. Specify the duration (in days) beyond which **Detailed** reports cannot be generated, against the **DetailedTime** parameter, and save the **eg_services.ini file**. For instance, to make sure that **Detailed** reports are disallowed for a **Timeline** of over 2 weeks, set the **DetailedTime** parameter to 14 and save the file. Say, subsequently, you attempt to generate a **Detailed** report for a **Fixed Timeline** of 3 weeks (which is greater than the 14-day limit set in our example). The instant you select the 3 weeks option from the **Fixed** list box, the **Detailed** option gets automatically disabled, and the **Trend** option gets enabled. Similarly, if you specify an **Any Timeline** that runs over 14 days, then, upon clicking the **SUBMIT** button to generate the report, a message box appears informing you that only the **Trend** option is permitted.

To proceed with the **Trend** report generation, click on the **OK** button in the message box. To terminate **Trend** report generation, click the **CANCEL** button.

- **What type of virtualization reports does eG Enterprise offer?**

  - **Physical Servers Report:** The **Physical Servers** report provides deep insights into the performance of a chosen physical server or a server farm during a specified timeline, and enables administrators to figure out the following:
    - Is the server adequately sized in terms of resources?
    - Which physical server in the farm has been consistently consuming resources excessively?
    - Which physical server in the farm hosts the maximum number of VMs?
    - Which physical server is the busiest in the farm in terms of disk activity and network traffic handled over time?
    - Which physical server in the farm has been the most problem-prone in the given time slot? What were the top problems encountered by that physical server?
  - **Virtual Machines - Overview:** The **Virtual Machines - Overview** report provides administrator with a closer look at the status and resource usage patterns of the VMs on a particular server or a server farm. Using this report, administrators can figure out the following:
    - How many VMs across the environment are in a powered-off state? Which ones are they?
    - What percentage of time has a VM been in a powered-off state?
    - Are VMs adequately sized in terms of resources?
    - Is any VM utilizing the allocated resources extensively?
    - How mildly/badly is the resource usage of a VM impacting the physical resources available with the server?

- **Virtual Machines - Details:** The **Virtual Machines - Details** report enables administrators to zoom into a particular VM, analyze its performance up, close, and accurately determine the following:
    - Has the VM been frequently moved to other hosts?
    - Which are the servers to which a VM has been Vmotioned during a given period?
    - Why was the VM moved? Is it because the VM is resource-intensive and is hence eroding the physical resources of the host, or is it owing to the fact that the host is not adequately sized to support the VM's operations?
    - How well is the VM using the resources allocated to it? Are spikes in usage sporadic/consistent? Have these usage patterns impacted the performance of any of the hosts on which the VM was Vmotioned/XenMotioned to during the designated period? What could be the reason for this anomaly - a resource-hungry application executing on the VM or insufficient resource allocation to the VM by the host?

- **VM Distribution Report:** In large virtualized farms, it is often difficult for administrators to know the VMs configured on the various virtual hosts and the current status of each VM. The use of VMotion/XenMotion technologies to migrate resource-intensive VMs to hosts that are well-sized, compounds the problem, as tracking the movement of VMs across hosts becomes a herculean task.

  The **VM Distribution** report is ideal for such farms as it provides administrators with the following:
    - a quick look at the VMs registered with a few/all virtual hosts in the environment;
    - the powered-on status of the VMs;
    - which VMs were VMotioned/XenMotioned to or from chosen virtual hosts in the environment during the configured timeline;

- **Virtual Centers Report:** This report tracks the performance history of vCenter installations in the target environment, so that administrators are able to perform the following effectively:
    - Determine how healthy the vCenter server as a whole has been over a broad period of time;
    - Take a closer look at the overall performance of the critical components of vCenter - i.e., clusters and individual physical servers managed by vCenter - over time;
    - Detect peculiar problem patterns, and analyze the root-cause of recurring performance issues.

- **Cluster Details Report:** The report provides you with historical insights into how well a chosen cluster / all clusters on a vCenter are utilizing the available resources. This way, resource-intensive clusters can be accurately identified. Also, by allowing administrators the flexibility to include the performance information related to the physical servers in the cluster, and the virtual machines executing on these physical servers, eG Enterprise enables administrators to zoom into the physical server / virtual machine that could be the root-cause of any resource contention at the cluster-level.

- **Cluster - Resource Pools Report:** As resource-intensive VM resource pools configured on one/more vSphere/ESX servers in a cluster can drain the physical resources of the corresponding virtual hosts, and consequently affect the aggregated resource base of the cluster, it is imperative that you monitor the current and historical performance of the cluster resource pools, so that resource contentions can be captured early and resolved quickly. The **Cluster - Resource Pools** report enables you to perform just that. Besides helping you identify resource-hungry resource pools in a cluster, this report also enables you to isolate the VM in the pool that is responsible for the abnormal behavior of the pool (if any), and the physical servers on which these VMs operate.

- **Cluster - Physical Servers Report:** While the **Cluster - Details** report enables you to identify resource-intensive clusters configured on a vCenter server, using the **Cluster - Physical Servers** report you can figure out which VM on which physical server in the cluster is responsible for the resource drain and why. In short, you can have the following questions answered with the help of this report:

    - Which physical servers are part of a cluster, and how have they been using the resources available to the cluster? Is any physical server using up excessive resources?
    - Are any VMs on the problem physical server cosumingg the allocated resources excessively?

- **Cluster - Virtual Machines Report:** A resource-intensive VM running on a host operating system can create a severe dent in not only the resources of the host, but also the consolidated resources of the cluster to which the host belongs. Ot is imperative that you monitor the current and historical performance of the VMs in a cluster, so that resource contentions can be captured early and resolved quickly. The **Cluster - Virtual Machines** report provides you with historical insights into the performance of VMs with a cluster, thereby turning the spotlight on potential resource contentions.

- **Datacenters Report**: The **Datacenters** report provides an overview of the performance of a chosen datacenter over time, besides allowing users the flexibility to focus on the overall performance of the clusters, physical servers, VMs, and datastores contained in a datacenter. From the report, administrators can understand the composition of the selected datacenter, and also analyze the overall availability and resource-effectiveness of the elements (physical servers/VMs/clusters/datastores) included in the datacenter.

- **VMotion Report**: Using the graphical/tabulated results provided by this report, administrators can perform the following:

- Easily observe and effectively analyze the migration activity over time, and understand the types of migration that were performed during the stated time;
- Detect failed migrations, identify the user(s) responsible for the same, and diagnose the reasons for the failure;
- Track the movement of specific VMs and accurately identify their current destination;
- Instantly identify under-performing hosts and/or resource-intensive guests in the target environment, on the basis of the migration activity ;

- **Datastore Report**: One of the key challenges faced by administrators of virtualized environments is making sure that the datastores do not run out of space at any point in time. This is because, if a datastore experiences a severe space crunch, the VMs associated with that datastore may fail to operate, and this would render the VMs inaccessible to users.

  To provide these administrators with historical insights into datastore usage so that they can capture potential space contentions, remove them, and ensure that business is transacted as usual in the virtualized environment at all times, eG Reporter offers the **Datastores Report**. Using this report, the following can be ascertained:

  - Are any datastores running out of space? If so, which ones are they?
  - Which VMs will be impacted by a space crunch in the datastore?

- **Uptime Report**: One of the key challenges faced by administrators of virtualized environments is making sure that the datastores do not run out of space at any point in time. This is because, if a datastore experiences a severe space crunch, the VMs associated with that datastore may fail to operate, and this would render the VMs inaccessible to users

➢ **What are the different types of Thin Client reports that are on offer?**

The eG Reporter provides the following Thin Client reports:

- The **Zone report**
- The **User report**
- The **Session report**
- The **Application report**

For elaborate details on how to generate these reports, refer to the *eG User Manual*.

➢ **Is the information displayed in a Thin client report configurable?**

Yes. Using specific sections of the **eg_report.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory), you can configure the data to be displayed in a thin client report. Please refer to the *eG User Manual* for more information.

- **Does the eG Reporter provide any Prediction Reports?**

  Yes, it does. **Prediction Reports** analyze the history of metrics and provide a prediction for the future. A **Prediction Report** applies built-in prediction mechanisms and sophisticated forecasting techniques on historical data for a metric to automatically compute how that metric is likely to change in the future. The result is a graph where historical data and future predictions are both plotted together.

  With the help of **Prediction Reports**, administrators can:

  - Clearly understand how load was and how it will change in the future;
  - Accurately determine when in the future the current resource capacity of the targets may get exhausted or may be inadequate; this information will enable you to re-evaluate your capacity plans and make sizing changes, so as to provide for such contingencies.
  - Isolate the exact times in the future when the overall performance of a target will dip, and also figure out how poorly the targets may perform during such times; potential threats to good health are thus revealed, and efforts to avert them can be initiated early.

  To generate the **Prediction Report**, select the **Prediction** option from the **Capacity Planning** menu in eG Reporter.

- **Does the eG manager have prediction capabilities?**

  There is a very minor difference between reaction and prediction. If the specific measure is undergoing a gradual change (for example the CPU utilization growing to near 100%), then the eG Enterprise Suite can easily detect this change and also display the trend values indicating that it has predicting capabilities. Alarms based on trends are not generated by the eG Enterprise Suite. However, there are some unpredictable events that occur due to external factors like disk failure, network failure etc.

- **What do the Cumulation and Correlation Reports help me achieve? How do I generate them?**

  The **Cumulation Report** plots the cumulative values of a chosen measure across descriptors/components, for a specified timeline. For instance, you can generate a **Cumulation Report** that plots the cumulative values of disk space usage for each of the disk partitions of a component. Likewise, you can configure a **Cumulation Report** to indicate the cumulative space usage of all servers that share the storage resources in a SAN environment. Such reports not only reveal the total space usage of the individual servers/disk partitions (as the case may be), but also reveal the overall space usage across the servers/disk paritions. This way, you will not only be able to pinpoint the resource-hungry elements in your infrastructure, but will also be able to assess how well/poorly the resource in its entirety was utilized during the given timeline. Based on these inputs, you can accurately forecast the future usage of your critical systems/servers, and plan their capacity accordingly. To generate this report, follow the *Capacity Planning -> Cumulation* menu sequence.

  In order to effectively plan the capacity of their IT environments, administrators have to closely observe how resource consumption / responsiveness / overall performance of a component varies vis-a-vis its load. For instance, administrators might want to check how disk space usage is affected by high disk activity or correlate the response time of a Web server with the connections to that server. The results of this correlation analysis will enable administrators to judge how

usage/response time will be affected by the anticipated load and accordingly ascertain the future capacity requirements of their critical servers/systems. The **Correlation Report** enables this correlation analysis. With the help of this report, administrators can pick two closely related measures and plot their historical values in a single graph, with one measure being the X axis and the other being the Y axis. For instance, you can pit disk space usage against free space on the disk in the same graph. To generate this report, follow the *Capacity Planning -> Correlation* menu sequence.

- **Does the eG Enterprise Suite have (basic) correction capabilities?**

  Yes. The eG Enterprise Suite includes a powerful Remote Control capability that allows monitor/supermonitor users to remotely manage and control servers from a web browser itself. From the browser, a supermonitor can execute commands on a monitored server, run some diagnosis, and initiate corrective actions to fix any problems.

- **Will Remote Control Actions cause security breaches in the target environment?**

  The control actions are enabled with no change in the eG architecture. The agents do not listen on any TCP ports. Hence, security risks in the target environment are minimum. Furthermore, since control actions can be initiated from a web browser, they can be triggered from any where, at any time.

- **What kind of commands can be issued to an agent host, remotely?**

  All non interactive commands that can be executed from command prompt (WIN) or $ prompt(UNIX), can be issued. For eg., On Unix platforms, commands for rebooting the system, script file execution, viewing files, viewing the processes running currently etc., can be issued; On Windows environments, commands for starting and stopping services / processes, viewing files, batch file execution etc., can be issued.

- **Does eG Enterprise provide a repository of ready-to-use commands for remote execution? If so, how can I access this command list?**

  Yes. By default, eG Enterprise provides a set of pre-defined remote commands, for use during remote problem correction or diagnosis. This list will be available to you as options of the **Command** list box in the **REMOTE CONTROL** page of the eG monitor interface.

- **Can I change this default command list?**

  Yes. If need be, administrators can override this default list using the **Manage Commands** page that appears upon following the Agents -> Settings -> Remote Control menu sequence.

- **Can I use the eG manager in addition to my existing SNMP manager?**

  To monitor applications, the eG manager can be integrated with existing SNMP managers such as OpenView's Network Node Manager (NNM) such that alarms generated by the eG Enterprise Suite can be prioritized and reported via NNM's alarm console. The combination of the eG Enterprise Suite and NNM can offer customers a powerful monitoring solution for networks, systems, and applications.

- **Can a customer update his/her own profile?**

  The administrator can modify a user's profile via the administrative interface of the eG manager. To reduce the burden on the central administrator, the eG Enterprise Suite also provides limited

functionality for monitor users to modify their profiles. On the eG Enterprise Suite's monitor interface, a monitor user can choose the **Profile** option of the menu running across the top of the eG Enterprise Suite's customer interface to modify his/her profile. From this page, the user can modify his/her password and alarm preferences.

- **Can I change the skin of the Admin, Monitor or Reporter consoles?**

  Yes, you can.

- **What are the skin colors available?**

  Blue, Grey, Green and Purple are available.

- **What is the default skin color for Admin, Monitor and Reporter?**

  The default Skin for the **Admin** console is **Blue**, **Monitor** console is **Gray** and for the **Reporter** console is **Purple.**

- **Can a customer choose the language in which the eG monitoring console should display data when he/she logs in?**

  Yes. A customer can set a language preference using the **USER PROFILE** page that appears upon clicking the **Profile** option in the menu running across the top of the eG Enterprise suite's customer interface.

- **What kind of additional privileges that customers enjoy by using the eG Enterprise Suite?**

  The eG Enterprise Suite provides limited functionality for monitor users to modify their profiles and to add new transactions to be monitored.

- **Explain the Quick Insight view**

  Often administrators may wish to track the values of certain key metrics in a single dashboard, so that they can proactively determine when the IT infrastructure may need attention. The metrics to be tracked may differ from one administrator to another. The eG Quick Insight offers an easy way for administrators to quickly track key metrics at each of the infrastructure tiers. With the help of this option, users can define infrastructure tiers to be monitored, add critical components for monitoring within each tier, and associate key metrics requiring closer observation with every component. Besides providing a holistic view of the environment at a single glance, this option enables users to focus on the more sensitive and critical components in the environment, and keep a close watch over their performance.

- **What are the key ingredients of a custom view?**

  - **Tiers**: A custom view is divided into broad sections known as tiers. A tier can be based on geography or it can match the different infrastructure tiers (for eg., the database tier, the web server tier), or it can even be a service / segment.

  - **Servers**: Each tier comprises of one or more **servers**. Typically, the critical servers in the monitored infrastructure are added to a tier.

  - **Metrics**: These are the key performance statistics that need to be extracted from every component in a tier.

- **What is the maximum number of tiers/servers/measures that can be defined for a custom Quick Insight view?**

  A maximum of 15 tiers can be configured. In every tier, a maximum of 15 servers can be added, and every server can be associated with an upper limit of 15 measures.

- **Can I add multiple servers to a tier, at one shot?**

  Yes. Once a new tier is added, the button alongside the tier name in the **DESIGN VIEW**. Clicking on this button will lead you to a page using which multiple servers can be simultaneously added to the tier. Similarly, a button appears alongside every server in a tier allowing you to add multiple measures to a server, at one shot.

- **Can I share the Quick Insight views I create with other users?**

  Yes, you can.

- **How do I share custom views with other users?**

  While configuring the **LAYOUT** of your view, you can pick an option from the **Sharing** list box, to indicate whether/not you want to share you view with other users.

- **What options are available in the Sharing list box?**

  - **Private**: If you set a view as **Private**, then all other users to the eG Enterprise system will be denied access to that view. Only the creator of the view will be able to access the view.

  - **Public**: If you set the view as **Public**, then only users with the following rights will have access to that custom view:

    - Users with access to all the managed components in the environment

    - Users with access to one/more components that are included in the view being shared

  - **Share:** If you choose the **Share** option from the **Sharing** list, then you can pick and choose the users who need to be granted access to that view

- **Can I hide views shared by other users? If so, how?**

  Yes, you can hide the unnecessary shared views. To execute this do the following.

  - Select the views you want to hide by providing checkmark in the corresponding checkboxes in the **QUICK INSIGHT VIEWS** page.

  - Then click on **Hide Selected Views**. Doing so ensures that the chosen views are hidden.

- **What is the purpose of a Live Graph Display?**

  Administrators of large, mission-critical environments are expected to be on high vigil 24 x 7, as even seemingly minor aberrations in performance could prove to be fatal for the infrastructure. Particularly, critical components of such infrastructures demand continuous attention. Therefore, it is essential for the monitoring solution in use to report even the smallest of variations in performance of such components, in real-time. Towards this end, the eG monitoring suite provides the **LIVE GRAPH DISPLAY** option which displays real-time graphs of key performance metrics relating to critical components in the infrastructure, allowing the administrator to keep a constant watch on the measure behavior, observe variations in the measures as they occur, and detect anomalies on-the-fly. Besides, the eG Enterprise Suite provides the option to plot historic data alongside the current data in the graph, so that an effective comparison of the past and the present performance can be performed, and appropriate performance decisions can be easily taken.

- **Mention the default number of rows and columns that a live graph view supports. What do they imply and what are the factors that I need to consider before modifying the default values in the live graph configuration page?**

  The **LIVE GRAPH DISPLAY** page will display a measure graph for every measure chosen in the live graph configuration page. These graphs will be displayed in a tabular format, characterized by rows and columns. Every graph will occupy a particular cell in the table. By default, the graph display will contain 2 rows and 2 columns. Accordingly, by default, the value **2** will be selected from the **No. of Rows** and the **No. of Columns** list boxes. This default layout can accommodate a maximum of 4 (2 rows x 2 columns) measure graphs, and all the 4 graphs can be viewed at a single glance. You can modify the layout by selecting a different number from the **No. of Rows** and the **No. of Columns** list boxes. While modifying the layout (i.e. the number of rows and columns), remember that to view all the graphs at a single glance **without using the scrollbar**, the layout for the graph display should at least provide for a cell each for each of the configured measures. For example, if 9 measures have been configured, then setting the **No. of Rows** to 3 and the **No. of Columns** to 3 will ensure that you do not have to scroll down or scroll right to view the graphs. The eG interface will automatically resize the graphs, so that all the 9 measure graphs can be viewed at a single glance. On the other hand, if the number of measures configured is greater than the maximum number of measure graphs that can fit into the defined layout, then the additional measure graphs can be viewed only by **scrolling**. The nature of this scrolling can be set by selecting the **Vertical** or the **Horizontal** option. Selecting the **Vertical** option will ensure that additional graphs are appended to the bottom of the graph display page. Therefore, you will have to scroll down to view the graphs. To scroll right to view the graphs, select the **Horizontal** option.

- **Explain the significance of the 'Interval' and 'No. of Times' fields in the live graph configuration page.**

  If the **Lookback** option is chosen, then proceed to select an **Interval** for the past values. For example, if the **Timeline** for the current measures is set to 1 hour, then the current measure graph will be plotted for the last one hour - say, for 12.00 PM to 1.00 PM of February 22, 2005. If the **Interval** is set to 1 day, then the eG Enterprise Suite will plot the values reported between the same hour (i.e., 12.00 PM and 1.00 PM), but for the previous day – i.e., February 21, 2005. To view meaningful data, it is recommended that the **Interval** set for the past values be greater than or equal to the **Timeline** chosen for the current values.

  The number of past measurements to be plotted can be specified in the **No. of Times** text box. In the example above, if **2** is specified against the **No. of Times** text box, then the measure graph will plot 2 sets of past measurements. This includes:

- The values for 12.00 PM to 1.00 PM on the February 21, 2005
- The values for 12.00 PM to 1.00 PM on the February 20, 2005

➢ **Is the 'Sharing' feature available only for Quick Insight?**

No. Live graph views, Custom dashboards, My Dashboards, and reports saved as favorites (using eG Reporter) can also be shared with other users, in the same manner discussed for quick insight views.

➢ **In a huge Citrix/Terminal/virtual desktop environment, if I want to locate a particular user and the servers accessed by him, which option in the eG monitor console should I choose?**

The **Users Search** option in the **Options** menu of the eG monitor console provides you with the information you are looking for.

Large Citrix/Terminal server farms and virtualized environments are typically characterized by a wide range of servers, each of which could be servicing a large number of concurrent users. When a particular user reports an issue with a server, the thin client administrator has to locate that user and the server that is being accessed by him/her, so that effective prognosis can be performed and the root-cause of the issue identified. Administrators of these vast environments often have to toil for hours to identify the problem servers. The **Users Search** option provided by the eG Enterprise suite is tailored for such large environments. This option lists the servers accessed by a specified user and the last access time for each server, so that administrators can quickly and accurately identify the server that was being accessed by the user when he/she reported a problem. Moreover, to support the administrator's diagnostic efforts, this page provides links to tools such as graphs and detailed measures, thus facilitating swift problem isolation and resolution.

➢ **How does VM Search help the users?**

Large virtualized environments typically comprise of hundreds of host operating systems supporting heterogenous virtualization technologies (VMware, Xen, LDOMs, etc.), with tens of virtual machines configured on each host. When the users in such environments report that they are unable to access a particular VM, administrators often take hours to figure out which server the VM is operating on and what its current state is. To enable administrators to quickly determine the location of a VM and its powered-on state, eG Enterprise offers the **VM Search** capability.

➢ **How to access the VM Search page?**

- Select **VM Search** from the **Options** menu.
- Once the **VM SEARCH** page appears, specify the name of the VM to search for in the **Virtual Machine Name** text box. To search for all VMs with names that begin with a particular string, just enter the string in the **Virtual Machine Name** text box.
- To know whether the VMs with the specified name are in a powered on state or not, choose **Powered on** from the **State** list. To view the complete list of VMs with the specified name, regardless of state, choose the **Any** option from the **State** list.
- The default **Timeline** for the VM search is **24 hours**. You can however, change this **Timeline**.
- To view the VMs that fulfill the specified conditions, click the **Submit** button.

- **What are the details displayed in the VM Search?**

  - The name of the VM
  - The server on which the VM is currently executing
  - The server type
  - The time at which the VM information was reported by the eG agent

- **What is the graph icon and diagnosis icon meant for in the VM Search page?**

  - By default, clicking on the **Graph** icon corresponding to a VM , reveals a graph of the powered-on state and physical CPU usage of that VM for the **Timeline** specified. This graph brings to light CPU usage excesses by the VM, and also intermittent changes in the state of the VM.
  - Similarly, clicking on the **Diagnosis** icon corresponding to a VM, by default, lists all the VMs that are executing on the **Server** displayed, the IP address of each VM, and the operating system.

Chapter

8

Configuration Management

- **Why should component/device configuration be monitored and managed?**

  Today's IT implementations are increasingly more complex. A single server can contain thousands of configuration elements, including system files, kernel parameters, registry keys, application settings, and firmware switches. Each of these elements may need to meet specific IT business requirements. Since a typical organization may have hundreds or even thousands of servers, the number of configurations to be tracked and managed can reach millions parameters.

  Change within an IT environment is a constant and can range from the planned, such as application and operating system upgrades, patch installations, and approved configuration updates, to the unplanned, including accidental system alterations and malicious security breaches. It is often these "unplanned" changes that have the greatest impact on the organization. Leading analyst firms estimate that, on average, more than 60 percent of all critical system and application outages are caused by inappropriate changes.

  The costs associated with unplanned and uncontrolled change can impact an organization on many levels. Unplanned and uncontrolled change can lead to a longer time-to-value for new products and services, and can cause inconsistent and unpredictable service. Uncontrolled change can also increase security and compliance risks, opening an infrastructure to malicious attacks and limiting an organization's ability to apply compliance strategies or the principles of good corporate governance. Auditors who discover evidence of uncontrolled change are likely to cite the organization for deficiencies or, in the most severe case, a material weakness.

  Finally, uncontrolled change can increase administrative costs as systems fail to provide the level of service expected and IT teams are forced to duplicate efforts and repeat processes in an effort to ensure systems remain functional.

- **What are the challenges in monitoring and managing the current configuration and the changes to it?**

  Traditional methods of managing and monitoring configuration settings are impeded by IT staffs that simply do not have time or resources to look at each element of a complex infrastructure individually. This can initially result in systems being deployed into the infrastructure that are not fully configured to a defined standard. Even when deployed properly, over time, the lack of visibility into the environment will result in "configuration drift" - configuration settings that, without staff knowledge, have changed over time until they are far from what they are supposed to be. Configuration drift negatively impacts an enterprise's operational performance and availability, security, and, eventually state of compliance to internal and external standards.

- **How doe eG Enterprise help in managing configuration?**

  What administrators require is a single, integrated solution that can collect, consolidate, and present in a central interface, the basic configuration and configuration change information related to all the components in the environment. eG Enterprise is such a solution. The optional **Configuration Management** module offered by this solution employs agent-based and agentless mechanisms to extract critical configuration and change details from each of the managed components in the environment, stores the data so collected in a central repository, and allows administrators to periodically query on the data via a 100%, web-based, easy-to-use **Configuration Management** console so that, the following tasks can be performed with elan:

  - From time-to-time, take stock of the applications, operating systems, devices, software, hardware, and services that are available in the environment;
  - Quickly access the basic configuration information pertaining to any system/application in the environment;
  - Accurately identify systems on which critical services have stopped, or on which mandatory software is missing;
  - Detect unplanned/unauthorized configuration changes with minimal effort;
  - Assess how a configuration change could have influenced overall performance/health of the system/application;
  - Run periodic checks to verify whether the entire infrastructure adheres to defined standards or not, and thus isolate deviations;

- **Is the eG Configuration Management capability license-controlled?**

  Yes, it is.

- **How does the eG Configuration Management capability function?**

  eG Enterprise determines the current configuration of a component/device and the changes to the configuration in about the same way it determines the level of performance achieved by that component/device - in other words, by executing tests on the target components. Like **Performance** tests, the eG agents execute **Configuration** tests on the targets to extract the configuration details of the target. These tests can be executed in an agent-based or in an agentless manner.

- **Where are the Configuration metrics displayed?**

  You can view the current configuration of a component/device and swiftly detect changes to it using the 100%, web-based **Configuration Management** console that eG Enterprise provides. Using this console, the following questions can be answered quickly and accurately:

  - Which are the platforms on which a particular application is currently running?
  - What are the applications currently executing on a particular system?
  - Which versions of an operating system are currently available in the target environment? Do any of these operating systems require an upgradation?

- What is the current configuration of an application?
- Are all mandatory software (like antivirus software) available and running on all managed systems? Which system does not have such software?
- Are all services critical to the functioning of your system and applications, up and running?
- Has the hotfix/patch been applied to all target systems?
- Do any systems require additional hard disk space? If so, which ones?
- Do any systems require their RAM size to be increased? If so, which ones?
- Have all Windows systems been updated with their latest service pack?
- Which are the systems that are supported by a particular processor family?
- Which systems in your environment have been assigned static IP addresses, and which ones hold dynamic IP addresses?
- How many printers has a system been configured with? What is the current status of each printer?
- Have any configuration changes occurred during the last 24-hours? If so, when exactly, and what are its details? Was this a planned change or an accidental one? Could this configuration change have induced a drop in the performance level of the system/application?
- Can any difference be noticed in the configuration of two components of the same type? If so, could this difference be the cause for the poor performance of one of the components?
- Which are the systems that fulfill a specific configuration requirement?

For more details about the console, please refer to the *eG Configuration Management* console.

Chapter 9

Conclusions

This manual clarifies common questions that users may have regarding the eG product. For more information on the eG Enterprise Suite of products, please visit our web site at **www.eginnovations.com**. For more details regarding the eG architecture and the details of the metrics collected by the eG agents, refer to the following documents:

- *A Virtual, Private Monitoring Solution for Multi-Domain IT Infrastructures*
- *The eG User Manual*
- *The eG Installation Guide*
- *The eG Measurements Manual*

Please contact us at support@eginnovations.com for further clarifications.